Approximation and Even Arity Hardness for Boolean Constraint Satisfaction

Zhijie Zheng[1] Yucong He[2] Chengzhi Xie[3]

[1]Peking University
[2]Beijing Institute of Technology
[3]National University of Singapore

Abstract

For every fixed $k \geq 5$, we give a polynomial-time approximation with ratio $(k+1)/2^k - \varepsilon$ for general Boolean constraint satisfaction of arity at most k. This improves the recent $k/2^k$ guarantee for $k \geq 10$ and matches the known odd-arity hardness bound. Consequently, the optimal deterministic polynomial-time ratio is $(k+1)/2^k$ for every odd $k \geq 5$ under the Unique Games Conjecture and P $\neq$ NP, and for $k = 2^r - 1 \geq 7$ under P $\neq$ NP alone, with arbitrarily small slack.

For every even $k \geq 4$, we also strictly lower the Unique Games hardness bound $(k+2)/2^k$ by an explicit rational amount. The hardness holds for unweighted conjunctions on exactly k distinct variables. Together, the two results narrow the interval of possible approximation ratios for every even $k \geq 6$.

Our semidefinite programming algorithm uses a reference sign shared by all variables to turn each conjunction into two antipodal Gaussian orthants in dimension $k+1$. Bakshi's variational bound, based on Mulgund's centered product inequality, reduces the analysis to a scalar estimate that we prove in every dimension at least six. The randomized implementation has polynomial dependence on $\log(1/\varepsilon)$; the deterministic implementation has polynomial dependence on $1/\varepsilon$.

The hardness construction uses two oppositely correlated hubs to obtain a negative Gaussian stability estimate for every measurable odd function taking values in $[-1, 1]$, uniformly in the dimension. This estimate allows a mixture of two entire tests to increase completeness linearly in the mixing weight while increasing soundness only quadratically. A common-function invariance argument transfers the bound to a Unique Games reduction.

Contents

1 Introduction **1**
1.1 Results . . . 1
1.2 The algorithm: one shared sign supplies the extra dimension . . . 3
1.3 Even arity hardness: negative stability and a mixture of tests . . . 3
1.4 Relation to prior work . . . 4
1.5 Organization and notation . . . 5

2 Predicates and normalization **5**

3 A common reference semidefinite relaxation **6**

4 The Gaussian variational inequality **7**

5 A scalar estimate in every dimension at least six **8**
5.1 Parametrization and the maximizing tilt . . . 8
5.2 The interval below one . . . 9
5.3 The interval above one reduces to its endpoints . . . 10
5.4 An elementary endpoint certificate . . . 10

6 One reference sign and two orthants **11**

7 Finite-bit implementation **12**
7.1 A rational interior point and feasible optimization . . . 13
7.2 Approximating the global score factor . . . 14
7.3 A finite Gaussian sampler and its coupling . . . 15
7.4 Sign stability and the approximation guarantee . . . 16

8 Deterministic rounding with polynomial accuracy dependence **17**
8.1 Rational coefficients with an independent noise floor . . . 17
8.2 Constant-dimensional Gaussian integration . . . 18
8.3 Fixing the shared Gaussian coordinates . . . 19
8.4 Fixing the remaining signs and accounting for losses . . . 20

9 An exact finite gap instance **21**

10 Known hardness and complexity assumptions **21**

11 Feasible local laws for even arity **22**
11.1 A mean-canceling two-hub law . . . 23
11.2 A positive-correlation law . . . 24

12 Dimension-free Gaussian negative stability **24**
12.1 Hermite and Fock notation . . . 24
12.2 The exact two-hub identity . . . 26
12.3 A global test mixture and regularization . . . 27

13 A common-function Boolean to Gaussian transfer **28**
13.1 Full support gives dimension-independent smoothing 29
13.2 One odd polynomial and its clipping 29
13.3 A vector Lindeberg comparison with explicit error 29
13.4 Applying the Gaussian bound to one common function 30

14 The Unique Games reduction **30**
14.1 Folding and the verifier 31
14.2 Completeness 31
14.3 Neighborhood averaging gives a common function 31
14.4 Influence decoding 32
14.5 Distinct variables and exact rational multiplicities 32
14.6 Conclusion of the hardness proof 33

15 Discussion **33**

A An exact arithmetic certificate **34**

1 Introduction

Boolean Max-k-CSP asks for an assignment maximizing the total weight of satisfied Boolean predicates, each on at most k variables. The predicates may vary from one constraint to another. A guarantee for this general problem applies to every Boolean predicate family of arity at most k, so its sharp approximation ratio describes the worst-case loss imposed by the arity alone.

Semidefinite programming gives a general description of this threshold. Under the Unique Games Conjecture (UGC), Raghavendra [Rag08] characterizes the optimal approximation of each fixed CSP through a canonical relaxation and a generic rounding scheme. Explicit numerical ratios still require an analysis of the relevant rounding and hardness constructions. For general Boolean Max-k-CSP, algorithms achieve order $k/2^k$ [CMM09, MM14]. Recent Gaussian inequalities sharpen this guarantee: Bakshi [Bak26] obtains $k/2^k$ for $k \geq 10$, while Makarychev [Mak26] obtains $(1 - o_k(1))k/2^k$ asymptotically, in both cases up to arbitrarily small algorithmic slack. On the hardness side, for $k \geq 3$, De and Mossel [DM13] give the Unique Games bounds $(k+1)/2^k$ for odd k and $(k+2)/2^k$ for even k. The resulting interval leaves room to improve both the algorithmic guarantee and the hardness bound at fixed arities.

We give the ratio $(k+1)/2^k - \varepsilon$ for every fixed $k \geq 5$. The numerator $k+1$ matches the known odd-arity hardness bound: for every odd $k \geq 5$, it gives the optimal deterministic polynomial-time threshold under UGC and P $\neq$ NP. Chan's hardness theorem [Cha12] gives the same conclusion under P $\neq$ NP alone when $k = 2^r - 1 \geq 7$. For even arity, we also strengthen the hardness bound, strictly lowering $(k+2)/2^k$ for every even $k \geq 4$. Thus both sides of the interval improve for even $k \geq 6$, while its exact threshold remains open. All threshold statements allow arbitrarily small slack.

These improvements come from two ways of using Gaussian correlation. On the algorithmic side, a reference sign shared by all variables lets each conjunction use two antipodal orthants in one higher dimension. The analysis must preserve this gain for every clause in a single assignment and throughout the finite-dimensional range needed for $k \geq 5$. On the hardness side, a test with higher dictator acceptance can also increase the acceptance of competing functions. We construct a test whose negative stability bounds that increase uniformly over all common odd decision functions taking values in $[-1, 1]$, then combine it with a test of higher dictator acceptance. The resulting improvement is quantitatively small and strict at every even arity from four onward.

1.1 Results

Write

$$\alpha_k = \frac{k+1}{2^k}. \tag{1}$$

All input weights and requested accuracy parameters below are rational. The input length includes the representation of the accuracy parameter; equivalently, one may supply a dyadic accuracy bound.

Theorem 1.1 (Approximation algorithm)**.** *For every fixed integer $k \geq 5$ and every rational $\varepsilon > 0$, there is a randomized algorithm for Boolean Max-k-CSP with nonnegative rational constraint weights that returns an exactly Boolean assignment x satisfying*

$$\mathbb{E}\,\mathrm{val}(x) \geq (\alpha_k - \varepsilon)\,\mathrm{OPT}\,.$$

The running time and number of random bits are polynomial in the explicit input bit length and in the requested precision $\log(1/\varepsilon)$ when $0 < \varepsilon < 1$. The guarantee is in expectation over the algorithm's finite random bits.

The guarantee applies to arbitrary incidence patterns and arbitrary optimum value. For fixed k, expanding truth tables into conjunctions increases the instance size by a factor depending only on k. The exact-arithmetic rounding attains α_k; approximate optimization and finite Gaussian sampling account for the numerical slack.

Theorem 1.2 (Deterministic approximation)**.** *For every fixed integer $k \geq 5$ and every rational $\varepsilon > 0$, a deterministic algorithm returns an exactly Boolean assignment x with*

$$\mathrm{val}(x) \geq (\alpha_k - \varepsilon)\,\mathrm{OPT}\,.$$

Its running time is polynomial in the explicit input bit length and $1/\varepsilon$, with an exponent that may depend on k.

This implementation uses conditional expectations and constant-dimensional Gaussian integration. An independent scalar-noise component keeps every conditional integration problem nonsingular.

Theorem 1.3 (Exact relaxation gap)**.** *For every $k \geq 5$, the common-reference relaxation in Section 3 satisfies*

$$\inf_{\mathcal{I}:\mathrm{OPT}(\mathcal{I})>0} \frac{\mathrm{OPT}(\mathcal{I})}{\mathrm{SDP}(\mathcal{I})} = \alpha_k.$$

The infimum is attained by the instance on k variables containing all 2^k signed k-literal conjunctions with equal weight.

The finite gap instance certifies that the rounding analysis attains the best uniform guarantee relative to this relaxation. Known hardness results give a stronger conclusion at odd arities: the ratio equals the UGC hardness ratio of De and Mossel [DM13, Theorems 5 and 16], and Chan's theorem [Cha12, Theorem 1.1 and Appendix E] gives matching unconditional NP-hardness for $k = 2^r - 1 \geq 7$. Thus Theorem 1.2 reaches the optimal deterministic approximation threshold in these cases, with the corresponding complexity assumptions. Section 10 states the promise gaps and explains the stronger assumption needed to exclude improvements by randomized algorithms with expected-value guarantees.

For even k, De and Mossel's hardness bound is $(k+2)/2^k$. Our second main result strictly improves it.

Theorem 1.4 (Strict even arity hardness improvement)**.** *Fix an even integer $k \geq 4$, and define the positive rational constants*

$$\begin{aligned} &n = k-2, && t = \frac{1}{16n^2}, && c = \frac{n(n-1)}{4}, && C = 2^{k-1}-1, \\ &d = (k-1)t, && a = \frac{cd}{C^2}, && x = ad, && \xi = \frac{x}{8(1+x)}. \end{aligned} \tag{2}$$

There are explicit rational completeness and soundness constants γ, β satisfying

$$\gamma \geq \frac{1+7x/8}{k+2}, \qquad \beta \leq 2^{-k}(1+x/2),$$

such that, assuming the Unique Games Conjecture, for every sufficiently small $\varepsilon > 0$ it is NP-hard to distinguish Max-k-AND instances of value at least $\gamma - \varepsilon$ from those of value at most $\beta + \varepsilon$. Constraints may be required to have exactly k distinct variables, and instances may be unweighted. In particular the hardness bound for general Boolean Max-k-CSP is at most

$$R_k = \frac{k+2}{2^k}\frac{1+x/2}{1+7x/8} = \frac{k+2}{2^k}\left(1 - \frac{3x}{8+7x}\right) < \frac{k+2}{2^k}. \tag{3}$$

A hardness bound R_k excludes every fixed deterministic polynomial-time approximation factor greater than R_k, under UGC and P $\neq$ NP. The reduction gives the possibly stronger ratio β/γ. The construction has completeness close to $1/(k+2)$ and gives a small explicit improvement through a dimension-independent negative-stability mechanism. For even $k \geq 6$, the remaining interval is $[\alpha_k, R_k]$, in the usual arbitrarily small slack sense. At $k = 4$, Theorem 1.4 supplies the hardness improvement; the algorithmic results apply from $k = 5$ onward.

1.2 The algorithm: one shared sign supplies the extra dimension

A value-preserving expansion reduces arbitrary predicates to signed conjunctions. The rounding analysis then relates each clause's satisfaction probability to its SDP value while producing one consistent assignment. Our relaxation ties every clause vector to a common reference direction: it assigns unit vectors $e, d_1, \ldots, d_n$ and a vector z_C to each conjunction. Its clause equations force

$$\langle e, z_C \rangle = \langle \sigma_{C,i} d_i, z_C \rangle = \|z_C\|^2 = q_C$$

for every literal $i \in C$. Equivalently, the two label vectors of each variable are orthogonal and sum to the same reference vector e.

The rounding forms noisy Gaussian scores $S_0, S_1, \ldots, S_n$ from these vectors and returns $x_i = \operatorname{sgn}(S_i)\operatorname{sgn}(S_0)$. A clause is satisfied exactly when its k oriented variable scores agree with the reference score. It therefore succeeds on two antipodal orthants of a $(k+1)$-dimensional Gaussian. All clauses use the same reference score and therefore refer to a single consistent assignment.

The clause equations give a common covariance component of size q_C. After adding independent noise, the covariance R satisfies

$$R \succeq \left(\frac{\pi q_C}{m} + \frac{1 - \pi/m}{m} \right) J, \qquad m = k + 1.$$

Bakshi's variational inequality [Bak26], derived from Mulgund's centered product inequality [Mul26], reduces the required orthant bound to a one-variable optimization. We prove that bound for every $m \geq 6$. We parametrize the optimizer by the inverse Mills ratio. A derivative argument reduces the required estimate to one endpoint, where an exact rational certificate completes the proof. Each orthant then contributes at least $mq_C/2^m$, giving $\alpha_k q_C$ for the two-orthant event.

The equal-weight instance containing every sign pattern shows that both the clause estimate and the whole relaxation analysis are exact. For the finite-bit implementation, the algorithm computes a rational factor for the joint Gaussian distribution and couples finite quantile samples to its ideal Gaussian input. A union bound on score-sign changes controls the objective loss. The deterministic implementation fixes shared coordinates while retaining independent scalar noise, so only constant-dimensional Gaussian probabilities need be evaluated.

1.3 Even arity hardness: negative stability and a mixture of tests

De and Mossel's even-arity test has dictator acceptance $1/(k+2)$, while pairwise independence makes its Gaussian acceptance 2^{-k} for every common odd decision function taking values in $[-1, 1]$. Our construction introduces a quantitative decrease in Gaussian acceptance that can offset the increase caused by a test with higher dictator acceptance.

We begin with their slice law. A mean-canceling perturbation preserves unbiased marginals and the all-positive atom, while creating two hubs: one is positively correlated with every remaining coordinate and the other is negatively correlated with those same coordinates. The remaining

coordinates are pairwise independent. We construct the Boolean law explicitly; the covariance therefore comes with a realizable local test.

For a common measurable odd function $f : \mathbb{R}^N \to [-1, 1]$, expand the conjunction acceptance over subsets of coordinates. Symmetry and cancellation leave only terms involving both hubs and an even number of leaves. In Hermite generating-function notation, these terms are $-\langle F^2, F^{2h}\rangle_t$. The leading term is $-2c\|F^2\|_t^2$. A Fock-space product inequality bounds every higher term relative to it, uniformly in the input dimension. Writing $S_t(f) = \sum_{j\geq 1} t^j W_j(f)$ for the noise stability, the resulting test has normalized acceptance at most $1 - cS_t(f)^2$.

A positive-correlation test raises the dictator completeness and has normalized Gaussian acceptance at most $1 + CS_t(f)$. Selecting one entire test with probability a and the other with probability $1 - a$ gives

$$1 - (1 - a)cu^2 + aCu, \qquad u = S_t(f).$$

Completeness grows linearly with the mixing weight, while maximizing this quadratic gives a soundness increase of quadratic order. Choosing explicit rational parameters produces the strict improvement in (3).

The mixture selects its branch once, then samples all columns conditionally independently from that branch. Its two branches use the same folded function. This property survives the Unique Games reduction: independent sampling of neighbors turns their tables into one neighborhood average. We prove smoothing and invariance using a common clipped polynomial, then decode an influential coordinate with bounds independent of the alphabet. Finally, cloning vertices and removing the small mass of colliding queries yields constraints on distinct variables, and rational multiplicities yield unweighted instances.

1.4 Relation to prior work

Trevisan's conjunction expansion [Tre98], also used in [MM14], preserves the value of every assignment when arbitrary predicates are replaced by weighted signed conjunctions. Hast [Has05] gave an $\Omega(k/(2^k \log k))$ approximation. Charikar, Makarychev, and Makarychev [CMM09] obtained a constant times $k/2^k$ approximation using semidefinite programming and Gaussian rounding; their treatment also uses Hast's reduction to AllEqual. Makarychev and Makarychev [MM14] refined this framework and developed the separate-label and clause-vector relaxation used in later analyses.

Mulgund [Mul26, Lemma 4.1 and Appendix B] proves a centered Gaussian product inequality for bounded log-concave functions. Bakshi [Bak26] uses it to obtain a Gaussian orthant variational bound and a $k/2^k$ approximation for $k \geq 10$. Makarychev [Mak26] independently obtains the asymptotic factor $(1 - o_k(1))k/2^k$ from Mulgund's comparison theorem. We apply Bakshi's full variational inequality, Equation (9) of the cited version, and prove the scalar estimate needed in dimensions $m \geq 6$.

Common reference vectors and relative-sign hyperplane decoding appear in Goemans and Williamson's algorithms for satisfiability [GW95, Section 7.2]. Our algorithm combines this structure with the clause-vector consistency equations and the Gaussian inequalities above. The separate-label relaxations [MM14, Bak26, Mak26] allow distinct label sums and round variable scores individually. Enforcing a common label sum makes each conjunction a two-orthant event in dimension $k + 1$; the scalar estimate for $m \geq 6$ gives the factor $k + 1$ for every fixed $k \geq 5$. The conjunction expansion, AllEqual viewpoint, Gaussian comparison theorems, and general SDP optimization machinery supply the underlying tools. Mahajan and Ramesh [MR99] derandomize SDP rounding through low-dimensional integration. Independent scalar noise keeps the conditional clause distributions nondegenerate. The resulting deterministic implementation has polynomial dependence on $1/\varepsilon$ for

fixed k. Our finite-bit analysis provides an affine feasibility repair and a quantile coupling tailored to this relaxation.

Austrin and Mossel [AM09] give a general UGC-based approximation-resistance criterion using balanced pairwise-independent distributions supported on satisfying assignments. De and Mossel [DM13] obtain the sharper odd- and even-arity bounds for Max-k-AND with smaller completeness. Our hardness proof builds on their slice distribution, folded long codes, and neighborhood-averaging and decoding framework. Global mixtures of local laws appear in the approximation-resistance frameworks of Austrin and Khot [AK13, Figure 3] and Khot, Tulsiani, and Worah [KTW13]. In those works, an even predicate is invariant under simultaneous sign reversal; here even arity refers to the number of variables. Austrin and Håstad [AH12, Theorem 8.6 and Lemma 10.3] use non-pairwise-independent tests and signed Gaussian stability of folded functions.

The Fock product estimate is the L^2 case of the Wick-product inequality of Da Pelo, Lanconelli, and Stan [DPLS11, Theorem 4.1], who credit the earlier L^2 estimate to Kuo, Saitô, and Stan [KSS02]. We include its coefficient proof to specify the normalization. Smoothing, hypercontractivity, and invariance are standard tools [Mos10]; our direct vector Lindeberg proof tracks one common function across both covariance branches. The new ingredients are the feasible mean-canceling two-hub law, its quantitative negative stability for every measurable odd function taking values in $[-1, 1]$, and its combination with the positive-correlation test.

1.5 Organization and notation

Sections 2–9 give the algorithm, scalar estimate, implementations, and exact relaxation gap. Section 10 states the known odd-arity and special-arity hardness results. Sections 11–14 prove the even-arity hardness theorem. Appendix A records the rational certificate.

For symmetric matrices, $A \succeq 0$ and $A \succ 0$ denote positive semidefiniteness and positive definiteness. A correlation matrix is positive semidefinite with unit diagonal. The symbols J and $\mathbf{1}$ denote the all-ones matrix and vector of the relevant dimension. Standard normal density and distribution functions are ϕ and Φ. Vector inequalities inside Gaussian probabilities are coordinatewise. Logarithms are natural.

2 Predicates and normalization

An instance consists of variables in $\{-1, 1\}$ and nonnegative rational weights on predicates of arity at most k. If a predicate lists the same variable several times, identify those positions and replace it by the induced truth table on its distinct variables. Thus arity means the number of distinct variables on which a constraint is evaluated. Zero-weight and identically false constraints may be discarded. Constant true predicates are allowed and are handled by the same construction, starting from arity zero.

We use the standard conjunction expansion [Tre98, MM14]. For a predicate P of arity $r \le k$ and weight w, introduce a conjunction for each accepting input $a \in \{-1, 1\}^r$, with the same weight w. For every assignment, exactly one such conjunction is true when P is true and none is true otherwise. This replacement preserves the objective pointwise.

For each resulting conjunction on $r < k$ variables, introduce $k - r$ private filler variables and replace it by all 2^{k-r} signed completions, each of weight w. Exactly one completion is true precisely when the original conjunction is true, independently of the filler values. Hence restriction to the original variables preserves the value of every completed assignment, and every original assignment has extensions with the same value. The optimum is unchanged. At most 2^k final conjunctions arise from any original predicate, so for fixed k the expanded instance has polynomial size.

Write a final conjunction as

$$C: \quad x_i = \sigma_{C,i} \quad (i \in C), \qquad |C| = k, \quad \sigma_{C,i} \in \{-1, 1\}.$$

Let w_C be its weight and let $W = \sum_C w_C$. If $W = 0$, any assignment is optimal. Otherwise divide every expanded weight by W, so that

$$\sum_C w_C = 1, \qquad 0 \le \mathrm{val}(x) \le 1, \qquad \mathrm{OPT} \ge 2^{-k}. \tag{4}$$

The final inequality follows because a uniformly random assignment satisfies each conjunction with probability 2^{-k}. The expansion may change the total weight. We normalize after expansion; undoing this positive scaling preserves every approximation ratio. The lower bound in (4) converts an additive numerical error into a multiplicative approximation guarantee.

3 A common reference semidefinite relaxation

Introduce one reference vector e, one decision vector d_i for each variable, and one vector z_C for each conjunction. The relaxation is

$$\begin{aligned} \text{maximize} \quad & \sum_C w_C q_C, && && (5)\\ \text{subject to} \quad & \langle e, e\rangle = 1, \qquad \langle d_i, d_i\rangle = 1 && \text{for every } i,\\ & q_C = \langle z_C, z_C\rangle = \langle e, z_C\rangle && \text{for every } C,\\ & \langle \sigma_{C,i} d_i, z_C\rangle = q_C && \text{for every } C \text{ and } i \in C, \end{aligned}$$

with all inner products represented by a single positive semidefinite Gram matrix. Thus every equality is affine in the matrix entries and the objective is linear. The rank is unrestricted. Positive semidefiniteness and Cauchy–Schwarz give

$$q_C \ge 0, \qquad q_C^2 = \langle e, z_C\rangle^2 \le \|e\|^2 \|z_C\|^2 = q_C,$$

so $q_C \in [0, 1]$ as a consequence of the affine equations and positive semidefiniteness.

Every Boolean assignment embeds feasibly: choose a unit vector e, set $d_i = x_i e$, and set $z_C = e$ if C is satisfied and $z_C = 0$ otherwise. The objective is exactly the assignment value. Consequently

$$\mathrm{SDP} \ge \mathrm{OPT}\,. \tag{6}$$

Every feasible vector has norm at most one, so the feasible Gram region is closed and bounded. In particular the optimum is attained.

The label-vector interpretation makes the common-reference condition explicit. Put

$$v_{i,+} = \frac{e + d_i}{2}, \qquad v_{i,-} = \frac{e - d_i}{2}. \tag{7}$$

Then $v_{i,+} \perp v_{i,-}$ and $v_{i,+} + v_{i,-} = e$ for every i. The required literal vector has inner product q_C with z_C and the opposite literal vector has inner product zero. Conversely, orthogonal label vectors summing to a common unit e recover unit decision vectors by subtraction. All variables therefore share the same sum of label vectors.

Remark 3.1. A relaxation encoding consistent full distributions on every constraint scope is stronger: it also imposes higher local moment identities. This distinction matters for the exact gap in Section 9.

4 The Gaussian variational inequality

We state Mulgund's centered product inequality and derive Bakshi's variational bound in the form needed by the rounding analysis.

Theorem 4.1 (Centered product inequality of Mulgund)**.** *Let R be an $m \times m$ correlation matrix and $X \sim N(0, R)$. Suppose each $f_i : \mathbb{R} \to [0, \infty)$ is bounded and log-concave, and for $Z \sim N(0, 1)$,*

$$\mathbb{E} f_i(Z) > 0, \qquad \mathbb{E}[Z f_i(Z)] = 0.$$

Then

$$\mathbb{E} \prod_{i=1}^{m} f_i(X_i) \geq \prod_{i=1}^{m} \mathbb{E} f_i(Z).$$

This is Lemma 4.1 of Mulgund [Mul26], whose Appendix B gives a direct proof. The statement applies to negative correlations and to asymmetric functions satisfying the centering condition.

For $b > 0$, define $a(b) > 0$ by

$$\int_{-\infty}^{b} x e^{a(b)x} \phi(x)\, dx = 0, \qquad F(b) = \log \int_{-\infty}^{b} e^{a(b)x} \phi(x)\, dx. \tag{8}$$

The tilt exists uniquely. Indeed, the integral on the left is strictly increasing in a, has value $-\phi(b) < 0$ at zero, and becomes positive for large a: the contribution of any fixed subinterval of $(0, b)$ diverges, whereas the negative contribution tends to zero. The implicit function theorem applies since the derivative with respect to a is a positive second-moment integral.

Proposition 4.2 (Bakshi's variational bound)**.** *If $R \succ 0$ is a correlation matrix and $X \sim N(0, R)$, then*

$$\log \Pr(X \geq 0) \geq \max_{b \in (0,\infty)^m} \left\{ \sum_{i=1}^{m} F(b_i) - \frac{1}{2} b^{\mathsf{T}} R^{-1} b \right\}. \tag{9}$$

Derivation following Bakshi [Bak26, Equation (9)]. Let $I(b) = e^{F(b)}$. Integration by parts, together with (8), gives

$$e^{a(b)b} \phi(b) = a(b) I(b), \qquad F'(b) = a(b). \tag{10}$$

When differentiating I, the chain-rule term involving $a(b)$ vanishes by centering. Differentiating the centering integral also gives $a'(b) < 0$: its partial derivative with respect to b is $b e^{a(b)b} \phi(b) > 0$ and its partial derivative with respect to a is positive.

The potential on the right of (9) attains its maximum in the interior, as the following boundary estimates show. Since the centered tilt minimizes $\int_{-\infty}^{b} e^{ax} \phi(x)\, dx$ over a, one has $I(b) \leq \Phi(b) < 1$, hence $F(b) < 0$. The negative definite quadratic term therefore forces the potential to tend to $-\infty$ as $\|b\| \to \infty$.

As $b \downarrow 0$, the centering equation forces $a(b) \to \infty$. Otherwise a bounded subsequence would converge to a nonnegative finite tilt with strictly negative first moment on $(-\infty, 0]$. A second integration by parts shows that the variance of the probability law proportional to $e^{a(b)x} \phi(x) \mathbf{1}_{x \leq b}$ is

$$1 - b a(b) > 0.$$

Thus $b a(b) < 1$, and (10) implies $I(b) \leq e\phi(0)/a(b) \to 0$. It follows that $F(b) \to -\infty$ at zero. These observations make every nonempty superlevel set of the potential compactly contained in $(0, \infty)^m$.

At a maximizer b^*, the gradient equation is

$$a(b^*) = R^{-1} b^*,$$

where a is applied coordinatewise. Set $a = a(b^*)$. The functions $f_i(x) = e^{a_i x}\mathbf{1}_{x \le b_i^*}$ satisfy all the hypotheses of Theorem 4.1. Completing the square in the multivariate Gaussian density gives

$$\begin{aligned}\Pr(X \le b^* - Ra) &= e^{-a^\mathsf{T} Ra/2}\mathbb{E}[e^{a^\mathsf{T} X}\mathbf{1}_{X \le b^*}] \\ &\ge \exp\left(\sum_i F(b_i^*) - \tfrac{1}{2}a^\mathsf{T} Ra\right).\end{aligned}$$

Since $Ra = b^*$, the event on the left is the negative orthant. Symmetry exchanges the negative and positive orthants, and $a^\mathsf{T} Ra = (b^*)^\mathsf{T} R^{-1} b^*$. □

Lemma 4.3 (Precision in the common direction)**.** *If $R \succ 0$ and $R \succeq \eta J$ for $\eta > 0$, then $\mathbf{1}^\mathsf{T} R^{-1}\mathbf{1} \le \eta^{-1}$.*

Proof. Multiplying on both sides by $R^{-1/2}$ gives $I \succeq \eta vv^\mathsf{T}$, where $v = R^{-1/2}\mathbf{1}$. The rank-one matrix on the right has nonzero eigenvalue $\eta\|v\|^2$, which is at most one. □

5 A scalar estimate in every dimension at least six

Lemma 5.1 (Uniform test point estimate)**.** *Let $m \ge 6$ be an integer, $q \in [0,1]$, and*

$$\delta = \frac{\pi}{m}, \qquad \eta = \delta q + \frac{1-\delta}{m}. \tag{11}$$

If $R \succ 0$ is an $m \times m$ correlation matrix satisfying $R \succeq \eta J$, then for $Y \sim N(0,R)$,

$$\Pr(Y \ge 0) \ge \frac{mq}{2^m}. \tag{12}$$

The proof optimizes the uniform test point in Proposition 4.2. We parametrize the optimizer, analyze the derivative on the interval below one, and reduce the interval above one to its endpoints. The case $q = 0$ is immediate, so assume $q > 0$ and put

$$t = mq \in (0, m], \qquad A = m\eta = 1 + \frac{\pi}{m}(t-1) > 0. \tag{13}$$

By Lemma 4.3 and the uniform vector $b\mathbf{1}$ in (9),

$$\log \Pr(Y \ge 0) \ge m \sup_{b>0}\left\{F(b) - \frac{b^2}{2A}\right\}. \tag{14}$$

5.1 Parametrization and the maximizing tilt

Define the inverse Mills ratio and a related positive function by

$$\lambda(z) = \frac{\phi(z)}{\Phi(z)}, \qquad h(z) = z + \lambda(z), \qquad \mu = \lambda(0) = \sqrt{\frac{2}{\pi}}.$$

The mean and variance of a normal variable conditioned to be at most z give

$$\lambda'(z) = -\lambda(z)h(z), \qquad h(z) > 0, \qquad 0 < r(z) := \lambda(z)h(z) < 1, \qquad h'(z) = 1 - r(z) > 0. \tag{15}$$

Direct integration gives the conditional mean $-\lambda(z)$ from the identity $x\phi(x) = -\phi'(x)$, and the second moment is $1 - z\lambda(z)$. Thus its variance is $1 - z\lambda(z) - \lambda(z)^2 = 1 - r(z) > 0$. Also $h(z)$ is the strictly positive conditional mean of $z - Z$, proving the positivity assertions.

Completing the square in (8) shows that at $b = h(z)$ the centering tilt is $a(b) = \lambda(z)$, and

$$F(h(z)) = \frac{\lambda(z)^2}{2} + \log \Phi(z). \tag{16}$$

The one-dimensional potential in (14) tends to $-\infty$ at both endpoints, as in the proof of Proposition 4.2. Its derivative is $a(b) - b/A$, which is strictly decreasing since $a'(b) < 0$. It therefore has a unique maximizer. The stationary equation in the above parametrization is

$$h(z) = A\lambda(z), \qquad z = \frac{\pi}{m}(t-1)\lambda(z). \tag{17}$$

There is exactly one such z. One can also see this directly from

$$\frac{d}{dz}\frac{z}{\lambda(z)} = \frac{1 + zh(z)}{\lambda(z)} = \frac{1 + h(z)^2 - r(z)}{\lambda(z)} > 0. \tag{18}$$

Existence follows from the maximization above; the derivative gives uniqueness and monotone dependence on t. At $t = 1$, the solution is $z = 0$; the sign of z equals the sign of $t - 1$.

Using $h = A\lambda$ in (16), the optimized value of the bracket in (14), after adding $\log 2$, is

$$L(z) = \log(2\Phi(z)) - \tfrac{1}{2} z\lambda(z).$$

It remains to prove

$$D_m(t) := mL(z(t)) - \log t \geq 0 \qquad (0 < t \leq m). \tag{19}$$

At $t = 1$ this quantity is zero.

5.2 The interval below one

Differentiating L and (18) yields

$$\frac{dL}{d(z/\lambda)} = \frac{\lambda^2}{2}, \qquad D_m'(t) = \frac{\pi}{2}\lambda(z)^2 - \frac{1}{t} = \frac{H_m(z)}{2t},$$

where

$$H_m(z) = \pi\lambda(z)^2 + mz\lambda(z) - 2. \tag{20}$$

Here we used $t = 1 + mz/(\pi\lambda(z))$. Also, $H_m(0) = 0$. Differentiating once more and substituting $r = \lambda h$ gives

$$\frac{H_m'(z)}{\lambda(z)} = m - mh(z)^2 - (2\pi - m)r(z). \tag{21}$$

For $z \leq 0$, monotonicity of h gives $h(z) \leq \mu$. The coefficient $2\pi - m$ in (21) changes sign between the two smallest permitted dimensions, so we treat $m = 6$ separately. If $m \geq 7$, then $m > 2\pi$ and the right side of (21) is at least $m(1 - 2/\pi) > 0$. If $m = 6$, use $r < 1$, $3 < \pi < 22/7$, and $2\pi - 6 > 0$ to obtain

$$\frac{H_6'(z)}{\lambda(z)} > 12 - \frac{12}{\pi} - 2\pi > 12 - 4 - \frac{44}{7} = \frac{12}{7} > 0.$$

Thus $H_m(z) < 0$ for $z < 0$. Therefore D_m decreases toward its value zero at $t = 1$, proving (19) on $(0, 1]$.

5.3 The interval above one reduces to its endpoints

For $z > 0$, define

$$I(z) = \int_0^z h(u)\,du, \qquad M(z) = \frac{2 - \pi\lambda(z)^2}{z\lambda(z)}.$$

Since $\lambda'/\lambda = -h$ and $\pi\mu^2 = 2$, one has

$$\lambda(z) = \mu e^{-I(z)}, \qquad M(z) = \frac{4\sinh I(z)}{\mu z}, \qquad \lim_{z\downarrow 0} M(z) = 4. \tag{22}$$

The derivative of M has the sign of

$$zh(z)\cosh I(z) - \sinh I(z).$$

This expression is positive: because h is increasing and positive, $zh(z) \geq I(z) > \tanh I(z)$. Thus M is strictly increasing. Moreover $M(z) \to \infty$, since $\lambda(z) \to 0$ and $z\lambda(z) \to 0$ as $z \to \infty$. Equation (20) becomes

$$H_m(z) = z\lambda(z)(m - M(z)).$$

For $m \geq 6$, the right-hand side is first positive and then negative on $(0, \infty)$, with exactly one change of sign. Consequently D_m on $[1, m]$ first increases and then may decrease; it has no interior minimum. Its minimum is attained at 1 or m. We already know $D_m(1) = 0$.

5.4 An elementary endpoint certificate

We prove $D_m(m) > 0$ for every $m \geq 6$. Evaluate the variational potential at the fixed test point corresponding to $z_0 = 1/2$. Set

$$a = \lambda(1/2), \qquad b = \tfrac{1}{2} + a, \qquad A_m = 1 + \pi - \frac{\pi}{m},$$

and define

$$G_m = \log(2\Phi(1/2)) + \frac{a^2}{2} - \frac{b^2}{2A_m}.$$

The optimization defining D_m implies

$$D_m(m) \geq mG_m - \log m. \tag{23}$$

As A_m increases with m, $G_m \geq G_6$. We will show $G_6 > 3/10$ using only rational estimates. This bound covers all these dimensions: $\log m/m$ decreases for $m \geq 6$, and the bound $\log 6 < 9/5$ proved below gives $\log 6/6 < 3/10$. We therefore need only certify that G_6 exceeds this threshold, using $z_0 = 1/2$ to lower-bound the optimized potential.

The classical bounds $25/8 < \pi < 22/7$ imply, by squaring,

$$\frac{199}{500} < \frac{1}{\sqrt{2\pi}} < \frac{2}{5}.$$

The degree-three and degree-four alternating Taylor sums at $1/8$ give

$$\frac{22}{25} < e^{-1/8} < \frac{353}{400}.$$

Multiplication yields

$$\frac{7}{20} < \phi(1/2) < \frac{353}{1000}. \tag{24}$$

Likewise $e^{-x^2/2} > 1 - x^2/2$ for $0 < x \le 1/2$, so

$$\Phi(1/2) > \frac{1}{2} + \frac{199}{500}\int_0^{1/2}(1 - x^2/2)\,dx = \frac{1}{2} + \frac{199}{500}\frac{23}{48} > \frac{69}{100}.$$

Since $\phi(x) < 2/5$, also $\Phi(1/2) < 7/10$. Hence

$$a > \frac{1}{2}, \qquad a < \frac{353}{690}, \qquad b < \frac{349}{345}, \qquad A_6 = 1 + \frac{5\pi}{6} > \frac{7}{2}. \tag{25}$$

For $x > 1$, put $u = (x-1)/(x+1)$. The positive expansion $\log x = 2\sum_{j\ge 0} u^{2j+1}/(2j+1)$ at $x = 69/50$ gives

$$\log(2\Phi(1/2)) > \log(69/50) > 2\left(\frac{19}{119} + \frac{1}{3}\left(\frac{19}{119}\right)^3\right) = \frac{1628072}{5055477}. \tag{26}$$

Combining (25) and (26),

$$\begin{aligned} G_6 &> \frac{1628072}{5055477} + \frac{1}{8} - \frac{121801}{833175} \\ &= \frac{482750044079}{1604608399800} = \frac{3}{10} + \frac{1367524139}{1604608399800} > \frac{3}{10}. \end{aligned} \tag{27}$$

Finally, $\log 6 < 9/5$, because the positive exponential Taylor sum through degree six is

$$e^{9/5} > \sum_{j=0}^{6} \frac{(9/5)^j}{j!} = \frac{7542629}{1250000} > 6.$$

The derivative of $(3/10)m - \log m$ is positive for real $m \ge 6$. Thus (23) implies

$$D_m(m) > \frac{3}{10}m - \log m \ge \frac{9}{5} - \log 6 > 0.$$

Together with the preceding two subsections this proves (19). Substituting it into (14) gives

$$\log \Pr(Y \ge 0) \ge -m\log 2 + \log(mq),$$

which completes the proof of Lemma 5.1. Appendix A gives an exact-arithmetic check of the displayed rational bounds.

6 One reference sign and two orthants

Let a feasible solution of (5) be fixed. Put $m = k+1$ and $\delta = \pi/m$. Since $k \ge 5$, $0 < \delta < 1$. Sample a standard Gaussian vector g in the joint vector space and independent standard Gaussian scalars $\xi_0, \xi_1, \ldots, \xi_n$. Form

$$S_0 = \sqrt{\delta}\,\langle g, e\rangle + \sqrt{1-\delta}\,\xi_0, \qquad S_i = \sqrt{\delta}\,\langle g, d_i\rangle + \sqrt{1-\delta}\,\xi_i. \tag{28}$$

Return the assignment

$$x_i = \operatorname{sgn}(S_i)\operatorname{sgn}(S_0). \tag{29}$$

There is exactly one reference score S_0, reused for all variables and all clauses. Every score has standard normal marginal, so ties occur with probability zero in the ideal rounding.

Lemma 6.1 (Clause probability)**.** *For every conjunction C and every feasible solution of* (5)*, rounding* (28)–(29) *satisfies*

$$\Pr(C \text{ is satisfied}) \geq \alpha_k q_C.$$

Proof. Write $q = q_C$. Orient the m relevant vectors as $w_0 = e$ and $w_j = \sigma_{C,i} d_i$ for the k variables of C. All are unit vectors, and the SDP constraints say

$$\langle w_j, z_C \rangle = q = \|z_C\|^2 \quad (0 \leq j \leq k).$$

Thus the residual vectors $u_j = w_j - z_C$ are orthogonal to z_C. Their Gram matrix H is positive semidefinite, its diagonal entries are $1 - q$, and the Gram matrix of the w_j is

$$\Sigma = qJ + H.$$

This formula remains valid at both $q = 0$ and $q = 1$, without dividing by $1 - q$.

Consider the oriented scores

$$Y = (S_0, (\sigma_{C,i} S_i)_{i \in C}).$$

The oriented noise variables remain independent standard Gaussians. The covariance of Y is the positive definite correlation matrix

$$R = \delta\Sigma + (1 - \delta)I.$$

With η as in (11),

$$R - \eta J = \delta H + (1 - \delta)(I - J/m) \succeq 0. \tag{30}$$

Also $R \succeq (1 - \delta)I \succ 0$. Lemma 5.1 therefore gives $\Pr(Y \geq 0) \geq mq/2^m$; if $q = 0$ the bound is simply zero.

By (29), C is satisfied exactly when all coordinates of Y have the same sign. These are the two disjoint open orthants $Y > 0$ and $Y < 0$. Gaussian symmetry gives them equal probability, and coordinate hyperplanes have probability zero. Hence

$$\Pr(C \text{ is satisfied}) = 2 \Pr(Y \geq 0) \geq \frac{2m}{2^m} q = \frac{k+1}{2^k} q.$$

□

Summing Lemma 6.1 with nonnegative weights gives, for any feasible SDP objective V,

$$\mathbb{E}\, \mathrm{val}(x) \geq \alpha_k V. \tag{31}$$

Linearity of expectation applies regardless of the dependence among clauses. At an exact optimum, (6) proves the ideal factor α_k. The next section turns this statement into Theorem 1.1.

7 Finite-bit implementation

We now implement the rounding on a randomized Turing machine. Throughout this section the expanded conjunction weights sum to one, M is the number of conjunctions, $N = n + 1$ is the number of shared score coordinates, $D = N + M$ is the size of the full SDP Gram matrix, and $m = k + 1 \geq 6$. The objective therefore lies in $[0, 1]$, and $\mathrm{OPT} \geq 2^{-k}$. The exact-arithmetic analysis applies to every feasible SDP matrix. We compute a rational feasible matrix whose objective is near optimal, then approximate the associated joint Gaussian rounding while controlling the probability of score-sign changes.

7.1 A rational interior point and feasible optimization

Let $\mathcal{A}$ be the affine space specified by the unit diagonals of $e, d_1, \dots, d_n$ and the equations

$$G_{e,z_C} = G_{z_C,z_C}, \qquad G_{d_i,z_C} = \sigma_{C,i} G_{z_C,z_C} \quad (i \in C).$$

Write $q_C = G_{z_C,z_C}$ and $f(G) = \sum_C w_C q_C$. The feasible region is $\mathcal{F} = \{G \in \mathcal{A} : G \succeq 0\}$. The bounds $0 \le q_C \le 1$ follow from its defining constraints: positive semidefiniteness gives $q_C \ge 0$ and the principal submatrix on e, z_C gives $q_C^2 \le q_C$. All diagonals, and consequently all entries, of every $G \in \mathcal{F}$ have absolute value at most one.

Lemma 7.1 (Quantitative rational interior point)**.** *The affine slice $\mathcal{F}$ contains an explicitly computable rational matrix G_0 with*

$$G_0 \succeq \lambda I_D, \qquad \lambda = \frac{1}{2M + 4m}.$$

Its entries have polynomial bit length. If $\mathcal{L} = \mathcal{A} - G_0$, then

$$G_0 + \{E \in \mathcal{L} : \|E\|_F \le \lambda/2\} \subseteq \mathcal{F} \subseteq G_0 + \{E \in \mathcal{L} : \|E\|_F \le 2D\}.$$

Proof. Set $t = 1/(2m)$ and $r = 1/(4m)$. Let $e, d_1, \dots, d_n$ and the private vectors f_C form an orthonormal family, and set

$$z_C = t\Big(e + \sum_{i \in C} \sigma_{C,i} d_i\Big) + \sqrt{r}\, f_C.$$

The squared norm of z_C is $mt^2 + r = t$, and every prescribed signed inner product with z_C is t. If C_0 is the N-by-M matrix whose Cth column is t times the corresponding signed incidence vector, the Gram matrix is

$$G_0 = \begin{pmatrix} I_N & C_0 \\ C_0^{\mathsf{T}} & C_0^{\mathsf{T}} C_0 + r I_M \end{pmatrix}.$$

This matrix is rational although the displayed vectors contain $\sqrt{r}$. Moreover, $\|C_0\|_{\mathrm{op}}^2 \le \|C_0\|_F^2 = M/(4m)$, and

$$\begin{aligned} (x,y)^{\mathsf{T}} G_0 (x,y) &= \|x + C_0 y\|_2^2 + r\|y\|_2^2, \\ \|x\|_2^2 + \|y\|_2^2 &\le 2\|x + C_0 y\|_2^2 + (2\|C_0\|_{\mathrm{op}}^2 + 1)\|y\|_2^2 \\ &\le (2M + 4m)(\|x + C_0 y\|_2^2 + r\|y\|_2^2). \end{aligned}$$

This proves the eigenvalue bound. A perturbation of Frobenius norm at most $\lambda/2$ has operator norm at most $\lambda/2$, proving the inner inclusion. The entry bound above gives $\|G\|_F \le D$ for every feasible G, including G_0, and proves the outer inclusion. □

The rational SDP optimization theorem of Grötschel–Lovász–Schrijver [GLS88], in the explicit formulation of de Klerk and Vallentin [dKV16, Theorem 1.1], now applies with inner radius $\lambda/2$ and outer radius $2D$. Consequently, for every rational $0 < \tau \le 1$, one can compute a rational matrix $G \in \mathcal{F}$ such that

$$f(G) \ge \mathrm{SDP} - \tau \tag{32}$$

in time polynomial in the explicit instance bit length and $\log(1/\tau)$. The output satisfies the affine equations exactly and is positive semidefinite. Lemma 7.1 supplies the theorem's rationality and quantitative-interior hypotheses.

A weakly feasible matrix can also be repaired explicitly. Suppose a symmetric rational X satisfies

$$\operatorname{dist}_F(X, \mathcal{F}) \le u \le 1, \qquad f(X) \ge \mathrm{SDP} - u.$$

Replace each shared-vector diagonal by 1 and each prescribed (e, z_C) or (d_i, z_C) entry by X_{z_C, z_C} or $\sigma_{C,i} X_{z_C, z_C}$, respectively, and update the symmetric entries accordingly. Leave every clause diagonal unchanged and call the result Y. Each entry changes by at most $2u$, so

$$Y \in \mathcal{A}, \qquad f(Y) = f(X), \qquad Y \succeq -\rho I_D, \qquad \rho = (2D+1)u.$$

Here $X \succeq -uI_D$ follows from its distance to $\mathcal{F}$ and $\|Y - X\|_{\mathrm{op}} \le 2Du$. Define the rational matrix

$$\widehat{G} = (1-\theta)Y + \theta G_0, \qquad \theta = \frac{\rho}{\lambda + \rho}. \tag{33}$$

It belongs to $\mathcal{A}$ and is positive semidefinite, since $-(1-\theta)\rho + \theta\lambda = 0$. Also $f(Y) \le 1 + u \le 2$ and $f(G_0) = t \ge 0$, hence

$$f(\widehat{G}) \ge \mathrm{SDP} - u - 2\theta \ge \mathrm{SDP} - u - 2\rho/\lambda.$$

Taking $u = \tau\lambda/[4(2D+1)]$ makes the last loss at most τ. All corrections and the mixing coefficient have polynomial bit length. The bounds on q_C again follow from positive semidefiniteness and the exact affine equations.

7.2 Approximating the global score factor

Let B be the N-by-N principal block of the feasible G indexed by $e, d_1, \ldots, d_n$. Thus $B \succeq 0$ and $B_{ii} = 1$. The entire ideal score vector has covariance

$$H = \delta B + (1-\delta) I_N, \qquad \delta = \pi/m. \tag{34}$$

In particular, if $A = H^{1/2}$ and $Z \sim N(0, I_N)$, then $S = AZ$ has exactly the joint law used by the rounding. This representation includes one reference coordinate shared by all clauses. Since $m \ge 6$ and $\pi < 4$,

$$\tfrac{1}{3} I_N \prec H \preceq N I_N, \qquad \|A_{i,*}\|_2 = 1. \tag{35}$$

The implementation factors this N-dimensional shared block; the full D-dimensional matrix supplies its feasibility certificate.

Lemma 7.2 (Rational score factor)**.** *For every rational $0 < \mu \le 1$, a rational matrix $\widetilde{A}$ with $\|\widetilde{A} - H^{1/2}\|_{\mathrm{op}} \le \mu$ can be computed in time polynomial in the bit length of B, N, and $\log(1/\mu)$.*

Proof. Compute a rational $p \in (3, 4)$ such that $|p - \pi| \le m\mu/(4N)$, and put $\delta' = p/m$ and $H' = \delta' B + (1-\delta') I_N$. Then $H' \succeq I_N/3$, $\operatorname{tr} H' = N$, and $\|H' - H\|_{\mathrm{op}} \le N|\delta' - \delta| \le \mu/4$. The two matrices are functions of B, so they have a common orthonormal eigenbasis. Applying $|\sqrt{a} - \sqrt{b}| = |a-b|/(\sqrt{a} + \sqrt{b})$ to their eigenvalues gives

$$\|H'^{1/2} - H^{1/2}\|_{\mathrm{op}} \le \mu/4.$$

The square root can be approximated explicitly in polynomial time. Put $X = I_N - H'/N$. Its spectrum lies in $[0, 1 - 1/(3N)]$. Write

$$\sqrt{1-x} = 1 - \sum_{j \ge 1} c_j x^j, \qquad c_j = \frac{\binom{2j}{j}}{4^j (2j-1)} \in (0, 1].$$

For $P_L = I_N - \sum_{j=1}^{L} c_j X^j$ we have $\|P_L\|_{\mathrm{op}} \le 1$ and

$$\|\sqrt{N} P_L - H'^{1/2}\|_{\mathrm{op}} \le 3N^{3/2} \exp(-(L+1)/(3N)).$$

Choose $L = 3N\lceil \log_2(12N^2/\mu)\rceil$ and approximate $\sqrt{N}$ by a rational number within $\mu/4$. The resulting rational multiple of P_L has the required accuracy. All coefficients, powers, and sums can be evaluated by rational arithmetic with polynomial bit length. For example, the required approximation to π follows from Machin's identity $\pi = 16\arctan(1/5) - 4\arctan(1/239)$ and the geometrically convergent arctangent series; scalar square roots can be evaluated by bisection. □

7.3 A finite Gaussian sampler and its coupling

Let Φ and ϕ denote the standard normal distribution function and density, and let $Q = \Phi^{-1}$. An explicit L^1 coupling controls the absolute sampling error.

Lemma 7.3 (Midpoint quantile sampling)**.** *For every integer $b \ge 2$, there is a sampler using exactly b independent unbiased bits and polynomially many bit operations in b which outputs a rational $\widetilde{Z}_b$. It admits a coupling to $Z \sim N(0,1)$ satisfying*

$$\mathbb{E}|Z - \widetilde{Z}_b| \le 20\sqrt{b}\, 2^{-b}. \tag{36}$$

Different calls can use independent copies of this coupling.

Proof. Let $h = 2^{-b}$. In the coupling take U uniform on $(0,1)$, put $J = \lfloor U/h \rfloor$, and define

$$Z = Q(U), \qquad Z_b = Q((J + \tfrac{1}{2})h).$$

The algorithm itself samples only the uniform b-bit integer J. It evaluates the indicated quantile to absolute error at most h and returns that rational approximation $\widetilde{Z}_b$.

We first bound the error of the exact midpoint-quantile sample. Set $a = -Q(h)$ and $a' = -Q(h/2)$. On each interior cell $[jh, (j+1)h]$, $1 \le j \le 2^b - 2$, monotonicity bounds the integral of $|Q(u) - Q((j+1/2)h)|$ by $h[Q((j+1)h) - Q(jh)]$. These differences telescope to $2ha$. On either end cell, symmetry and the substitution $u = \Phi(z)$ give the bound

$$\int_0^h |Q(u) - Q(h/2)|\, du \le \int_0^h |Q(u)|\, du + ha' = \phi(a) + ha'.$$

The elementary lower Mills bound $\overline{\Phi}(a) \ge a\phi(a)/(a^2+1)$ follows by integrating $\phi(z)/z$ by parts; it yields $\phi(a) \le h(a + a^{-1})$. Since $h \le 1/4$, we have $a > 1/2$. The Gaussian exponential-tail bound gives $a' \le \sqrt{2\log(2/h)} \le 2\sqrt{b}$. Consequently,

$$\mathbb{E}|Z - Z_b| \le h(4a + 2a^{-1} + 2a') \le h(6a' + 4) \le 16\sqrt{b}\, h.$$

Adding the quantile-evaluation error h proves (36).

The requested quantile can be computed in polynomial time by certified bisection. Every query probability lies in $[2^{-b-1}, 1 - 2^{-b-1}]$. Put $T = 2\lceil\sqrt{b+1}\rceil$. The exponential-tail bound brackets every requested quantile in $[-T, T]$, and throughout this interval

$$\phi(x) \ge 2^{-20(b+1)} =: \ell_b.$$

For example, $T^2 \le 16(b+1)$, $e < 4$, and $1/\sqrt{2\pi} > 1/3$ imply the stated lower bound. Compute Φ with certified absolute error at most $\kappa = \ell_b 2^{-b-4}$ during bisection. If its error interval lies entirely on

one side of the target probability, the corresponding half-interval is discarded. Otherwise the tested midpoint is within $2\kappa/\ell_b \le 2^{-b-3}$ of the desired quantile, by the lower bound on the derivative. After $O(b + \log b)$ comparisons, the interval width is at most 2^{-b}.

Finally, these certified CDF evaluations require only polynomial time. For rational $x \in [-T, T]$ use

$$\Phi(x) = \frac{1}{2} + \frac{1}{\sqrt{2\pi}} \sum_{j \ge 0} \frac{(-1)^j x^{2j+1}}{2^j j!(2j+1)}.$$

The integrated Taylor remainder after index L has absolute value at most $T(T^2/2)^{L+1}/(L+1)!$. Taking $L = O(b + p)$ makes it at most 2^{-p}; here the requested p is $O(b)$. Rational evaluation of this finite sum, together with a sufficiently accurate approximation to $1/\sqrt{2\pi}$, has polynomial bit complexity. The preceding arctangent and bisection procedures supply the latter constant. □

7.4 Sign stability and the approximation guarantee

Fix a target rounding loss $0 < \nu \le 1$. Couple N independent sampler outputs $\widetilde{Z}_1, \dots, \widetilde{Z}_N$ to independent standard normals $Z_1, \dots, Z_N$ as in Lemma 7.3. Set $S = AZ$ and compute $\widetilde{S} = \widetilde{A}\widetilde{Z}$ by exact rational arithmetic. If $\|\widetilde{A} - A\|_{\mathrm{op}} \le \mu \le 1$ and $\kappa_b = 20\sqrt{b}\,2^{-b}$, each coordinate obeys

$$\begin{aligned} \mathbb{E}|S_i - \widetilde{S}_i| &\le \mathbb{E}|((A - \widetilde{A})Z)_i| + \sum_j |\widetilde{A}_{ij}| \mathbb{E}|Z_j - \widetilde{Z}_j| \\ &\le \mu + \sqrt{N}(1+\mu)\kappa_b \le \mu + 2\sqrt{N}\,\kappa_b. \end{aligned} \tag{37}$$

The first term uses the variance of a linear form in standard normals; the second uses $\|A_{i,*}\|_2 = 1$ and Cauchy–Schwarz. To control all N score signs, we allocate an error probability of order ν/N to each coordinate. A threshold of order ν/N bounds the Gaussian mass near zero; Markov's inequality then requires an expected score error of order $(\nu/N)^2$. This determines the following choices. Take

$$\gamma = \frac{\nu^2}{16N^2}, \qquad \mu = \gamma/2, \qquad 20\sqrt{b}\,2^{-b} \le \frac{\gamma}{4\sqrt{N}}.$$

Such a b is $O(\log(N/\nu))$. It can be chosen using the stronger rational inequality $20b\,2^{-b} \le \gamma/(4N)$. Then (37) is at most γ. Each ideal coordinate S_i is standard normal, so, for any $h_0 > 0$,

$$\begin{aligned} \Pr[\mathrm{sign}(S_i) \ne \mathrm{sign}(\widetilde{S}_i)] &\le \Pr[|S_i| \le h_0] + \Pr[|S_i - \widetilde{S}_i| > h_0] \\ &\le \sqrt{2/\pi}\,h_0 + \gamma/h_0. \end{aligned} \tag{38}$$

Choose $h_0 = \nu/(4N)$. A union bound over all N coordinates makes the probability of any score-sign disagreement at most ν (in fact, at most $\nu/2$). When each computed sign agrees with its ideal counterpart, the two decoded Boolean assignments are identical, because both use the same reference coordinate. A computed score equal to zero is assigned a fixed tie-breaking sign. As the normalized objective is between zero and one, the expected objective loss from replacing the ideal assignment by the finite-bit assignment is at most ν.

For $0 < \varepsilon \le 1$, take

$$\tau = \nu = \frac{\varepsilon 2^{-k}}{2}.$$

For larger ε one may run the construction with $\min\{\varepsilon, 1\}$. Apply (32) and then the finite-bit rounding. Writing $\alpha_k = (k+1)/2^k$, the exact-arithmetic rounding inequality yields

$$\begin{aligned} \mathbb{E}[\mathrm{val}(\widetilde{x})] &\ge \alpha_k(\mathrm{SDP} - \tau) - \nu \\ &\ge \alpha_k\,\mathrm{OPT} - \varepsilon 2^{-k} \\ &\ge (\alpha_k - \varepsilon)\,\mathrm{OPT}\,. \end{aligned}$$

Here $\alpha_k \le 1$ and $\mathrm{OPT} \ge 2^{-k}$ justify the last two steps. For fixed k, the predicate expansion, SDP optimization, rational factor construction, and N sampler calls all have bit complexity polynomial in the explicit input length and $\log(1/\varepsilon)$. Only $Nb = O(N \log(N/\nu))$ random bits are used. The algorithm therefore returns a Boolean assignment with the claimed expected-value guarantee. Section 8 gives a deterministic construction with a different dependence on the precision.

8 Deterministic rounding with polynomial accuracy dependence

The randomized algorithm above has polynomial dependence on $\log(1/\varepsilon)$. Conditional expectations yield a deterministic algorithm with polynomial dependence on $1/\varepsilon$. The independent Gaussian noise is retained until all shared Gaussian coordinates have been fixed; this keeps every intermediate clause calculation nonsingular and constant dimensional.

Theorem 8.1 (Deterministic implementation)**.** *For every fixed integer $k \ge 5$ and rational $\varepsilon > 0$, there is a deterministic algorithm which, on an explicit nonnegatively weighted Boolean Max-k-CSP instance with rational weights, returns an assignment of value at least*

$$\left(\frac{k+1}{2^k} - \varepsilon\right) \mathrm{OPT}\,.$$

Its running time is polynomial in the total input bit length and $1/\varepsilon$, with an exponent depending on k.

First perform the same value-preserving expansion into conjunctions on k distinct variables. If no positive-weight conjunction remains, return any assignment. If $\varepsilon \ge \alpha_k = (k+1)/2^k$, any assignment also suffices. We may therefore assume $0 < \varepsilon < \alpha_k$, normalize the expanded weights to sum to one, and put

$$\eta = \varepsilon 2^{-k}/10. \tag{39}$$

In this nontrivial case $n \ge k$, $N = n + 1 \ge 1$, the objective is in $[0, 1]$, and $\mathrm{OPT} \ge 2^{-k}$.

8.1 Rational coefficients with an independent noise floor

Compute a rational feasible SDP matrix of value at least $\mathrm{SDP} - \eta$ by Section 7, and let B be its shared N-by-N Gram block. Its ideal score covariance is $H = \delta B + (1-\delta) I_N$, where $\delta = \pi/(k+1)$. For the present purpose use the following equivalent representation:

$$T = H - \tfrac{1}{4} I_N = \delta B + (\tfrac{3}{4} - \delta) I_N, \qquad S^* = T^{1/2} g + \tfrac{1}{2}\xi, \tag{40}$$

where $g, \xi \sim N(0, I_N)$ are independent. This score vector has covariance H and hence the same ideal joint rounding law, with the same single reference coordinate. Moreover, $T \succ I_N/12$ and $T_{jj} = 3/4$. The independent noise makes the factored matrix positive definite even when B is singular.

Put $\gamma = \eta^2/(16N^2)$. Choose rational $\delta' \in (0, 2/3)$ with $|\delta' - \delta| \le \gamma/(4N)$, and let

$$T' = \delta' B + (\tfrac{3}{4} - \delta') I_N.$$

Then $T' \succ I_N/12$, and $\|T' - T\|_{\mathrm{op}} \le \gamma/4$. Since T, T' commute, their common eigenbasis and the scalar square-root identity give

$$\|T'^{1/2} - T^{1/2}\|_{\mathrm{op}} \le \sqrt{3}\, \|T' - T\|_{\mathrm{op}} < \gamma/2.$$

The binomial construction in Lemma 7.2, with spectral floor $1/12$ in place of $1/3$, computes a rational U within $\gamma/2$ of $T'^{1/2}$ in operator norm. Indeed, scaling by N makes the residual spectrum lie in $[0, 1-1/(12N)]$, so $O(N\log(N/\gamma))$ terms suffice. All coefficient bit lengths and this preliminary computation are polynomial in the instance bit length and $\log(1/\eta)$.

Write u_j for row j of U and use the rational-coefficient scheme

$$S_j = u_j \cdot g + s\xi_j, \qquad s = \tfrac{1}{2}, \qquad j = 0, \dots, n. \tag{41}$$

We have $\|u_j\|_2 \le \sqrt{3/4} + \gamma \le 2$ and, in particular, $1/2 \le s \le 2$. Couple this scheme to (40) using the same g, ξ. Each score has expected absolute error at most γ. The anti-concentration calculation (38), applied to the ideal standard-normal marginals with threshold $\eta/(4N)$, bounds the probability of any score-sign disagreement by η. Thus the expectation of the rational-coefficient scheme is at least that of the ideal scheme minus η.

8.2 Constant-dimensional Gaussian integration

We need a deterministic numerical estimate of a clause probability after conditioning on shared coordinates. Conditioning generally introduces nonzero means.

Lemma 8.2 (Uniform integration at fixed dimension)**.** *Fix m. Given rational $a \in \mathbb{R}^m$, rational symmetric V with $V \succeq I_m/4$ and $V_{jj} \le 8$, and rational $0 < \tau < 1/2$, one can approximate the probability that $N(a, V)$ lies in a specified orthant, or in a specified pair of antipodal orthants, to additive error τ. The algorithm is deterministic and polynomial in the rational input bit length and $1/\tau$, with an exponent depending only on m. The error bound is uniform in a.*

Proof. Let $R = V^{1/2}$. We integrate in the standard Gaussian input coordinates: the event is defined by the signs of $a + RZ$, where $Z \sim N(0, I_m)$. Truncating these input coordinates makes the tail estimate independent of the mean a. The lower bound on the output variances will separately control the mass near every sign boundary, also uniformly in a. The rows of this square root have norms in $[1/2, \sqrt{8}]$, and $V \preceq 8mI_m$. Choose an integer $T = O(\sqrt{\log(m/\tau)})$ such that $2m\exp(-T^2/2) \le \tau/4$. For a standard normal vector $Z \in \mathbb{R}^m$,

$$\Pr[Z \notin [-T, T]^m] \le \tau/4.$$

Partition this box into congruent cubes of rational side length

$$h \le \frac{\tau}{64m^2};$$

one can subdivide each side into $\lceil 128Tm^2/\tau \rceil$ equal intervals. Compute a rational factor $\widetilde{R}$ with

$$\|\widetilde{R} - R\|_{\mathrm{op}} \le e := \frac{\tau}{64m^2(T+1)}.$$

The same binomial construction as above applies, now with eigenvalues in $[1/4, 8m]$, and has polynomial bit complexity. For a cube of center c, include its Gaussian mass if the rational vector $a + \widetilde{R}c$ belongs to the specified orthant or pair of orthants, with any fixed convention on boundary signs.

For any z in that cube, coordinate j satisfies

$$|(a + Rz)_j - (a + \widetilde{R}c)_j| \le \sqrt{8}\,\sqrt{m}\,h/2 + e\sqrt{m}\,T \le d := 2\sqrt{m}\,h + e\sqrt{m}\,T.$$

If the true and computed membership indicators differ, at least one coordinate has different signs, and hence $|(a+Rz)_j| \le d$ for some j. A normal variable of variance at least $1/4$ has density at most $2/\sqrt{2\pi}$, regardless of its mean. Consequently the Gaussian mass of all such misclassified points is at most

$$\sum_{j=1}^{m} \Pr[|(a+RZ)_j| \le d] \le 4md \le \frac{3\tau}{16\sqrt{m}} < \frac{\tau}{4}.$$

This bound controls both grid-boundary and factor-approximation errors uniformly in a. For two orthants the same bound applies: a change in membership in their union still requires a coordinate-sign change.

If the number of cubes is Q, the exact standard Gaussian mass of a cube $\prod_j [\ell_j, r_j]$ is

$$\prod_{j=1}^{m} (\Phi(r_j) - \Phi(\ell_j)).$$

Compute each CDF value with certified error at most $\tau/(8mQ)$ and clip each resulting interval probability to $[0,1]$. The telescoping bound for products in $[0,1]$ shows that the mass error per cube is at most $\tau/(4Q)$. Sum the selected rational masses exactly. The total error from these evaluations is at most $\tau/4$, so truncation, classification, and arithmetic together contribute less than τ. There are $Q = O((Tm^2/\tau)^m)$ cubes. The CDF computations follow the certified Taylor evaluation in Lemma 7.3, and need only polynomially many bits. The mean enters the rational sign tests; the grid size and error bound are uniform in the mean. □

8.3 Fixing the shared Gaussian coordinates

Keep all noises ξ_j unconditioned while fixing $g_1, \dots, g_N$ in order. After a rational prefix $g_1, \dots, g_{\ell-1}$ has been chosen, let $F_\ell(t)$ be the conditional expected objective given also $g_\ell = t$. For a clause C, orient its scores as $(S_0, (\sigma_{C,i} S_i)_{i \in C})$. Conditional on this prefix, this $m = k+1$ dimensional vector has a rational mean and rational covariance

$$V_C \succeq s^2 I_m = I_m/4, \qquad (V_C)_{jj} \le \|u_j\|_2^2 + s^2 \le 8.$$

The remaining common coordinates contribute a positive semidefinite covariance, while all the independent scalar noises remain present. The clause probability is

$$\Pr[Y_C \in [0,\infty)^m \mid \text{prefix}] + \Pr[Y_C \in (-\infty, 0]^m \mid \text{prefix}]. \tag{42}$$

Their overlap has probability zero. A nonzero conditional mean can give the two orthants different probabilities, so the computation retains both terms. Lemma 8.2 computes each clause's union probability. Since the clause weights sum to one, estimating every clause to error τ estimates $F_\ell(t)$ to error τ.

The functions F_ℓ are uniformly Lipschitz with constant $L = 4m$, independent of the prefix and N. To see this, couple the remaining randomness at two arguments t, t'. Coordinate j shifts by at most $2|t - t'|$. Conditional on all common Gaussian coordinates, the independent term $s\xi_j$ has density at most $1/(s\sqrt{2\pi}) \le 2/\sqrt{2\pi}$. Its sign therefore changes with probability at most $4|t - t'|/\sqrt{2\pi}$. A union bound over the m coordinates of a clause, followed by weighted summation, gives

$$|F_\ell(t) - F_\ell(t')| \le 4m|t - t'|.$$

Weighted summation preserves the bound for arbitrarily dependent clauses.

Choose an integer $T_0 = O(\sqrt{\log(N/\eta)})$ with $\Pr[|Z| > T_0] \le \eta/(8N)$ for $Z \sim N(0,1)$, and a rational grid $\mathcal{G} \subset [-T_0, T_0]$ including the endpoints and of mesh at most $\eta/(8NL)$. Map Z to a nearest grid point, clamping values outside the interval to an endpoint. Since $0 \le F_\ell \le 1$,

$$\max_{t \in \mathcal{G}} F_\ell(t) \ge \mathbb{E}[F_\ell(Z)] - L\frac{\eta}{8NL} - \frac{\eta}{8N} \ge \mathbb{E}[F_\ell(Z)] - \frac{\eta}{4N}.$$

Evaluate F_ℓ at every grid point to error $\eta/(8N)$ and choose an approximate maximizer. Its true value is within $\eta/(4N)$ of the true grid maximum. Thus fixing this coordinate loses at most $\eta/(2N)$ in conditional expectation. After all N coordinates are fixed, the total loss is at most $\eta/2$.

The grid contains $O(T_0 NL/\eta)$ points. At each point the integrations have dimension m, independently of the ambient dimension N. All chosen grid coordinates, conditional means, and conditional covariances are rational of polynomial bit length. For fixed k, this coordinate-fixing procedure takes time polynomial in the input length and $1/\eta$.

8.4 Fixing the remaining signs and accounting for losses

After g is fixed, the signs of the scores in (41) are independent. Put $p_i = \Phi((u_i \cdot g)/s)$. For a possible reference sign $r \in \{-1, 1\}$, define

$$V_r = \sum_C w_C \prod_{i \in C} \Pr[\operatorname{sign}(S_i) = r\sigma_{C,i}].$$

The current expectation is $p_0 V_+ + (1 - p_0)V_-$, so choosing the larger of V_+, V_- cannot decrease it. After fixing the reference, sequentially fix each variable score sign by maximizing its conditional expected objective. At every step the current expectation is a convex combination of the two candidate conditional objectives.

For an explicit numerical implementation, compute and clip each p_i, $1 \le i \le n$, to $[0, 1]$ with error at most

$$\zeta = \frac{\eta}{8kn}.$$

Evaluate all subsequent products and weighted sums exactly over the rationals. At any stage a conditional clause probability is either zero or a product of at most k remaining factors, each equal to a p_i or $1 - p_i$. The telescoping product bound gives error at most $k\zeta = \eta/(8n)$ for each conditional objective, including V_r. Thus the approximate reference choice loses at most $\eta/(4n) \le \eta/4$, and the n approximate variable choices lose at most $n \cdot \eta/(4n) = \eta/4$. Including the shared-coordinate stage, the total derandomization loss is at most η. The reference choice compares the two conditional objectives directly; p_0 only enters their convex combination. CDF evaluation for these rational arguments is polynomial-time as well: outside a threshold $O(\sqrt{\log(1/\zeta)})$, the Gaussian tail bound permits returning 0 or 1; inside it, the certified Taylor calculation applies. This evaluation has polynomial complexity for all rational means.

If the chosen score signs are $y_i \in \{-1, 1\}$, output $x_i = ry_i$. This is an exactly Boolean assignment and satisfies

$$\begin{aligned} \operatorname{val}(x) &\ge \alpha_k(\mathrm{SDP} - \eta) - \eta - \eta \\ &\ge \alpha_k \,\mathrm{OPT} - 3\eta \\ &\ge (\alpha_k - \varepsilon)\,\mathrm{OPT}, \end{aligned}$$

where the three losses are respectively SDP optimization, rational coefficient approximation, and conditional-expectation rounding. Equation (39) and $\mathrm{OPT} \ge 2^{-k}$ justify the final inequality. Normalization of positive rational weights increases bit length only polynomially, independently of the smallest positive weight. All deterministic searches and integrations are polynomial in the total input length, including the encoding of ε, and $1/\varepsilon$ for fixed k. This proves Theorem 8.1.

9 An exact finite gap instance

Take k variables and every signed conjunction

$$C_\sigma : \quad x_i = \sigma_i \ (1 \le i \le k), \qquad \sigma \in \{-1, 1\}^k,$$

each of weight one. Every assignment satisfies exactly one conjunction, so OPT $= 1$. Choose orthonormal vectors $e, d_1, \dots, d_k$, and for every sign vector put

$$z_\sigma = \frac{e + \sum_{i=1}^k \sigma_i d_i}{k+1}.$$

Then

$$\|z_\sigma\|^2 = \langle e, z_\sigma \rangle = \langle \sigma_i d_i, z_\sigma \rangle = \frac{1}{k+1}.$$

These choices satisfy all constraints simultaneously and give SDP objective $2^k/(k+1)$. All Gram entries are rational. There is also a direct matching upper bound. For any feasible solution, set $a_\sigma = e + \sum_i \sigma_i d_i$ and $m = k + 1$. The constraints and Cauchy–Schwarz give

$$m q_\sigma = \langle a_\sigma, z_\sigma \rangle \le \|a_\sigma\| \sqrt{q_\sigma}, \qquad q_\sigma \le \frac{\|a_\sigma\|^2}{m^2},$$

where the second inequality is immediate if $q_\sigma = 0$. Summing over all sign patterns cancels every cross term, so

$$\sum_\sigma q_\sigma \le \frac{1}{m^2} \sum_\sigma \|e + \sum_i \sigma_i d_i\|^2 = \frac{2^k}{m}.$$

Thus

$$\mathrm{SDP} = \frac{2^k}{k+1}, \qquad \frac{\mathrm{OPT}}{\mathrm{SDP}} = \alpha_k.$$

The lower bound OPT / SDP $\ge \alpha_k$ for all instances also follows from (31). This proves Theorem 1.3.

For this witness, the $k + 1$ base vectors are orthonormal, so the oriented scores have covariance I for every δ. Consequently each clause has probability exactly 2^{-k}, equal to $\alpha_k q_C$. Thus equality holds for each clause and for the complete instance.

For even k, this gap rules out a uniform guarantee of $(k + 2)/2^k$ relative to this relaxation's objective. Stronger relaxations that impose consistent local distributions can exclude this specific vector solution.

10 Known hardness and complexity assumptions

We state the existing hardness results as promise gaps. In this section values are fractions of total constraint weight in the original hardness instance; this convention is independent of the expanded normalization in Section 2.

Theorem 10.1 (De and Mossel [DM13, Theorem 16])**.** *For every fixed $k \ge 3$ and every sufficiently small $\zeta > 0$, the Unique Games Conjecture implies NP-hardness of distinguishing Max-k-AND instances with*

$$\mathrm{OPT} \ge \rho_k - \zeta \qquad \textit{from instances with} \qquad \mathrm{OPT} \le 2^{-k} + \zeta, \quad \rho_k = \frac{1}{2\lceil (k+1)/2 \rceil}.$$

The resulting limiting ratio is

$$\frac{2^{-k}}{\rho_k} = \frac{2\lceil (k+1)/2 \rceil}{2^k} = \begin{cases} (k+1)/2^k, & k \text{ odd}, \\ (k+2)/2^k, & k \text{ even}. \end{cases}$$

Conjunctions are a subclass of general Boolean predicates. Thus for odd $k \geq 5$, Theorem 1.1 attains this UG-hard ratio with arbitrarily small algorithmic slack. This is multiplicative hardness with completeness close to $1/(k+1)$.

Proposition 10.2 (Chan's special-arity consequence)**.** *If $k = 2^r - 1 \geq 7$, then for every sufficiently small $\zeta > 0$ it is NP-hard, without assuming UGC, to distinguish general Boolean Max-k-CSP instances with* $\mathrm{OPT} \geq 1 - \zeta$ *from instances with* $\mathrm{OPT} \leq \alpha_k + \zeta$.

Proof. Apply Chan's theorem [Cha12, Theorem 1.1] to the subgroup

$$\mathcal{C} = \left\{ (a \cdot v)_{v \in \mathbb{F}_2^r \setminus \{0\}} : a \in \mathbb{F}_2^r \right\} \subseteq \mathbb{F}_2^k.$$

It has size $2^r = k + 1$: the evaluation map is linear and injective, since it includes evaluation on every coordinate basis vector. For a uniform a, each nonzero v gives a uniform bit $a \cdot v$. Two distinct nonzero binary vectors are linearly independent, so each pair of distinct coordinates is uniform on $\mathbb{F}_2^2$. Hence $\mathcal{C}$ is balanced and pairwise independent. Chan's theorem gives completeness arbitrarily close to one and soundness arbitrarily close to $|\mathcal{C}|/2^k = (k+1)/2^k$. His Appendix E explicitly describes this binary simplex-code construction. The resulting additive constraints are Boolean predicates of arity k. □

Randomized algorithms and complexity assumptions. The preceding assertions and Theorem 1.4 give NP-hard promise gaps. Deterministic improvements are excluded under $P \neq NP$. For randomized expected-value guarantees, the relevant implication can be seen directly: let a promise gap have completeness c and soundness s, and suppose a randomized algorithm always outputs a feasible assignment and guarantees expected value at least $\beta\,\mathrm{OPT}$, where $\beta c > s$. In a completeness instance, if $V \in [0,1]$ is its output value, then

$$\Pr(V > s) \geq \frac{\mathbb{E}V - s}{1 - s} \geq \frac{\beta c - s}{1 - s} > 0.$$

In a soundness instance, every output has value at most s. For fixed gap parameters, independent repetition gives a one-sided randomized decision procedure for the promise problem. With the corresponding polynomial reductions, such a procedure for an NP-hard gap would imply $NP \subseteq RP$. Thus excluding randomized expected-value improvements uses the corresponding randomized complexity assumption or a randomized formulation of UGC. The deterministic interpretation uses $P \neq NP$.

Together with Theorem 1.2, these gaps identify α_k as the optimal deterministic polynomial-time approximation threshold, in the usual arbitrarily small slack sense, for odd $k \geq 5$ under UGC and $P \neq NP$, and for $k = 2^r - 1 \geq 7$ under $P \neq NP$ alone. The randomized algorithm attains the same ratio with a better dependence on precision.

11 Feasible local laws for even arity

Throughout Sections 11–14, $k \geq 4$ is even and the constants n, t, c, C, d, a, x, ξ are those in (2). In particular $n = k - 2 \geq 2$. These constants are independent of the instance size and the long-code

dimension. Since

$$C \geq k-1 = n+1, \qquad cd = \frac{n^2-1}{64n} < \frac{n}{64},$$

we have

$$0 < a < \frac{n}{64(n+1)^2} < \frac{1}{2}. \tag{43}$$

All probabilities below are rational. The laws have unbiased single-bit marginals and allow higher-moment asymmetry under simultaneous sign reversal.

11.1 A mean-canceling two-hub law

Let $X = (X_1, \ldots, X_k) \in \{-1, 1\}^k$ and $S = \sum_i X_i$. Start with the pairwise-independent even-arity law of De and Mossel [DM13, Fact 12]: it puts mass

$$q_0 = \frac{1}{k+2} \quad \text{on } S = k, \qquad A = \frac{1}{2} \quad \text{on } S = 0, \qquad B = \frac{k}{2(k+2)} \quad \text{on } S = -2, \tag{44}$$

and is uniform within each slice. Directly, $\mathbb{E}S = 0$ and $\mathbb{E}S^2 = k$, so exchangeability makes every mean and every off-diagonal pair moment zero. Set $h = -(k-1)t$ and alter the conditional densities on the two nontrivial slices to

$$1 + h(X_1 - X_2) \quad (S = -2), \qquad 1 - \frac{h(k-2)}{k}(X_1 - X_2) \quad (S = 0). \tag{45}$$

The conditional means of these densities are one, by exchanging coordinates 1 and 2. They are strictly positive: $2|h| = 2(n+1)/(16n^2) < 1$, and $(k-2)/k < 1$. The mass on the all-positive point remains q_0. Call the resulting probability law D_0.

For a uniform slice $S = s$ and distinct indices,

$$\mathbb{E}X_i = \frac{s}{k}, \qquad \mathbb{E}X_iX_j = \frac{s^2-k}{k(k-1)}, \tag{46}$$
$$\mathbb{E}X_iX_jX_\ell = \frac{s^3-(3k-2)s}{k(k-1)(k-2)}.$$

The first two formulas follow from S and S^2. In expanding S^3, the all-equal terms contribute S and the terms with exactly two equal indices contribute $3(k-1)S$, leaving $6\sum_{i<j<\ell} X_iX_jX_\ell$; this gives the third formula. Put $u = (1, -1, 0, \ldots, 0)$. Since $\sum_i u_i = 0$, a linear density tilt proportional to $\sum_i u_iX_i$ changes the mean of X_j by the proportionality constant times $(1 - \mathbb{E}X_iX_j)u_j$. The two weighted contributions to each mean cancel, because

$$A\left(-\frac{h(k-2)}{k}\right)\frac{k}{k-1} + Bh\frac{k^2-4}{k(k-1)} = 0.$$

For a pair of distinct coordinates i, j, the expectation of $X_iX_j\sum_\ell u_\ell X_\ell$ is

$$\left(\frac{s}{k} - \frac{s^3-(3k-2)s}{k(k-1)(k-2)}\right)(u_i + u_j).$$

It vanishes on $S = 0$. On $S = -2$ its weighted tilted contribution is

$$Bh\left(-\frac{2}{k} - \frac{6}{k(k-1)}\right)(u_i + u_j) = -\frac{h}{k-1}(u_i + u_j) = t(u_i + u_j).$$

Thus D_0 has zero means and covariance

$$\Sigma_0 = I + tA_*, \tag{47}$$

where A_* has $+1$ between hub 1 and each of the n leaves, -1 between hub 2 and each leaf, and zero on every other off-diagonal entry. Its diagonal is zero. The only nonzero eigenvalues of A_* are $\pm\sqrt{2n}$: on the span of the hub-difference vector and the leaf-sum vector it is the corresponding two-by-two off-diagonal matrix, and it vanishes on the orthogonal complement. Since $t\sqrt{2n} < 1$, Σ_0 is positive definite. The explicit Boolean law therefore realizes the covariance used in the Gaussian analysis.

11.2 A positive-correlation law

Let D_+ be uniform within the same three slices, now with masses

$$q_+ = \frac{1+(k-1)t}{k+2}, \qquad B_+ = \frac{kq_+}{2}, \qquad A_+ = 1 - \frac{(k+2)q_+}{2} = \frac{1-(k-1)t}{2} \tag{48}$$

on $S = k, -2, 0$, respectively. These are positive and sum to one. Moreover

$$\mathbb{E}S = kq_+ - 2B_+ = 0, \qquad \mathbb{E}S^2 = k(k+2)q_+ = k[1+(k-1)t].$$

Exchangeability and (46) show that every mean is zero and every off-diagonal covariance is t. Therefore

$$\Sigma_+ = I + t(J - I). \tag{49}$$

The all-positive atom has probability $q_+ = (1+d)/(k+2)$, strictly larger than q_0.

12 Dimension-free Gaussian negative stability

Let $f : \mathbb{R}^N \to [-1, 1]$ be any measurable odd function, in any finite dimension N. Oddness means $f(-z) = -f(z)$ almost everywhere. Throughout this section all Gaussian vector marginals are standard in $\mathbb{R}^N$. A single function f is used for all rows and in both branches.

12.1 Hermite and Fock notation

Let H_α be the probabilists' multivariate Hermite polynomial for a multi-index $\alpha \in \mathbb{N}^N$, with

$$\mathbb{E}H_\alpha(G)H_\beta(G) = \alpha!\mathbf{1}_{\alpha=\beta}.$$

Here $G \sim N(0, I_N)$, $|\alpha| = \sum_i \alpha_i$, and $\alpha! = \prod_i \alpha_i!$. Define the entire generating function

$$F(z) = \mathbb{E}\big[f(G)\exp(z\cdot G - z\cdot z/2)\big] = \sum_\alpha \frac{\mathbb{E}[f(G)H_\alpha(G)]}{\alpha!} z^\alpha. \tag{50}$$

The dot products in the exponential are bilinear, including for complex z. Boundedness of f makes the expectation locally dominated on complex compact sets, justifying its Taylor expansion. Its coefficients are real and only odd total degrees occur. For a power series U put

$$\|U\|_u^2 = \sum_\alpha \alpha! u^{|\alpha|} |[z^\alpha]U|^2, \qquad \langle U, V\rangle_u = \sum_\alpha \alpha! u^{|\alpha|} [z^\alpha]U\,\overline{[z^\alpha]V}.$$

These expressions are used only when the norms are finite. For F,

$$\|F\|_u^2 = \sum_{j\ge 1} u^j W_j(f) =: S_u(f), \qquad W_j(f) = \sum_{|\alpha|=j} \frac{\mathbb{E}[fH_\alpha]^2}{\alpha!}, \qquad \|F\|_1^2 = \mathbb{E}f^2 \le 1. \tag{51}$$

For $0 \le u \le 1$, it follows that $S_u(f) \le u$. If Z has independent standard circular complex Gaussian coordinates with $\mathbb{E}|Z_i|^2 = 1$, monomial orthogonality gives

$$\|U\|_u^2 = \mathbb{E}|U(\sqrt{u}Z)|^2. \tag{52}$$

For entire functions of finite norm this follows by integrating over coordinatewise phases and then radii, using the nonnegative coefficient sum. In particular it applies to all products used below once the next lemma gives finiteness.

Lemma 12.1 (Fock product estimate)**.** *For $p, q > 1$ with $p^{-1} + q^{-1} = 1$,*

$$\|UV\|_u \le \|U\|_{pu}\|V\|_{qu}.$$

Also $\|U^h\|_u \le \|U\|_{hu}^h$ for each integer $h \ge 1$.

Proof. This is the L^2 specialization of the Wick-product inequality [DPLS11, Theorem 4.1]; the following coefficient argument supplies the normalization. Write $U = \sum_\beta u_\beta z^\beta$ and $V = \sum_\gamma v_\gamma z^\gamma$. For each α, the weights

$$w_{\alpha,\beta} = \binom{\alpha}{\beta} p^{-|\beta|} q^{-|\alpha-\beta|} \quad (\beta \le \alpha)$$

sum to one. Weighted Cauchy–Schwarz gives

$$\begin{aligned} \alpha! u^{|\alpha|} |[z^\alpha](UV)|^2 &\le \sum_{\beta\le\alpha} \frac{\alpha! u^{|\alpha|} |u_\beta|^2 |v_{\alpha-\beta}|^2}{w_{\alpha,\beta}} \\ &= \sum_{\beta+\gamma=\alpha} \beta!(pu)^{|\beta|}|u_\beta|^2\, \gamma!(qu)^{|\gamma|}|v_\gamma|^2. \end{aligned}$$

Summing the nonnegative terms over α proves the first assertion. The same multinomial argument with h equal weights $1/h$ proves the second. For arbitrary series, each coefficient convolution is finite; summing the resulting nonnegative bound by Tonelli gives the same estimate whenever the right-hand side is finite. We use the Gaussian L^2 expansion below to evaluate bounded conditional expectations. □

The negative leading term in the two-hub expansion below is proportional to $\|F^2\|_t^2$. To control the higher terms for functions whose stability may be arbitrarily small, we retain one factor of $\|F^2\|_t$ in the bound on $\|F^{2h}\|_t$. Interpolation between two scales provides this relative estimate. Hölder's inequality on the nonnegative coefficient sum also shows that $\|U\|_u$ is log-convex in $\log u$. In particular, if $h \ge 2$ and

$$R = h^{h/(h-1)}t, \qquad ht = t^{1/h}R^{(h-1)/h},$$

then

$$\begin{aligned} \|F^{2h}\|_t &\le \|F^2\|_{ht}^h \le \|F^2\|_t \|F^2\|_R^{h-1} \\ &\le \|F^2\|_t (2R)^{h-1} = \|F^2\|_t\, 2^{h-1} h^h t^{h-1}, \end{aligned} \tag{53}$$

provided $2R \le 1$. For the last inequality, the product estimate and (51) give $\|F^2\|_R \le \|F\|_{2R}^2 = S_{2R}(f) \le 2R$. For $h \le n/2$ this condition holds: $h^{1/(h-1)} \le 2$, so $R \le 2ht \le nt = 1/(16n)$. The inequality $h^{1/(h-1)} \le 2$ follows from $h \le 2^{h-1}$ by induction. Finally, Jensen's inequality in (52) gives

$$\|F^2\|_t^2 = \mathbb{E}|F(\sqrt{t}Z)|^4 \ge \left(\mathbb{E}|F(\sqrt{t}Z)|^2\right)^2 = S_t(f)^2. \tag{54}$$

12.2 The exact two-hub identity

Let $(G_1, \dots, G_k)$ have covariance $(I + tA_*) \otimes I_N$ and define normalized conjunction acceptance

$$L_0(f) = \mathbb{E} \prod_{i=1}^{k} (1 + f(G_i)).$$

Consider its expansion over subsets of rows. Nonempty leaf-only subsets vanish because leaves are independent and $\mathbb{E}f = 0$. A subset with one hub and an even number of leaves has odd cardinality and vanishes by simultaneous sign symmetry. With one hub and an odd number of leaves, the two hub choices cancel: changing the sign of hub 2 changes all its $-t$ leaf correlations to $+t$, and changes the sign of f. The two hubs alone are independent and give zero. A subset with both hubs and an odd number of leaves again has odd cardinality and vanishes. The surviving subsets therefore contain both hubs and an even number $r \geq 2$ of leaves.

For one such subset replace G_2 by $-G_2$, writing the resulting hubs as U_1, U_2 and the selected leaves as $V_1, \dots, V_r$. The hubs are independent within their block, the leaves are independent within their block, and every cross-block covariance is tI_N. Put

$$H = \frac{U_1 + U_2}{\sqrt{2}}, \qquad L = \frac{V_1 + \cdots + V_r}{\sqrt{r}}, \qquad \rho = t\sqrt{2r} < 1.$$

Then $\mathrm{Cov}(H, L) = \rho I_N$. Every within-block orthogonal residual has zero covariance with both block sums and with the opposite block's residuals; joint Gaussianity gives the corresponding independence. For independent standard Gaussians $Z_1, \dots, Z_j$, let

$$A_j(z) = \mathbb{E}\left[\prod_{\ell=1}^{j} f(Z_\ell) \,\middle|\, \frac{\sum_\ell Z_\ell}{\sqrt{j}} = z \right].$$

This bounded measurable function is in L^2. Conditional expectation gives the unsigned mixed moment $\mathbb{E}[A_2(H)A_r(L)]$. To compute its Hermite coefficients, factor the exponential generating function across the independent Z_ℓ:

$$\mathbb{E}\left[\prod_{\ell=1}^{j} f(Z_\ell) e^{z\cdot(\sum_\ell Z_\ell)/\sqrt{j} - z\cdot z/2} \right] = F(z/\sqrt{j})^j.$$

Equating coefficients proves

$$\mathbb{E}[A_j(G)H_\alpha(G)] = \alpha! j^{-|\alpha|/2} [z^\alpha] F^j. \tag{55}$$

The L^2 Mehler expansion is absolutely convergent by Cauchy–Schwarz. Substitution in that expansion gives

$$\begin{aligned} \mathbb{E}[A_2(H)A_r(L)] &= \sum_\alpha \frac{\rho^{|\alpha|}}{\alpha!} \mathbb{E}[A_2 H_\alpha] \mathbb{E}[A_r H_\alpha] \\ &= \sum_\alpha \alpha! t^{|\alpha|} [z^\alpha] F^2 [z^\alpha] F^r = \langle F^2, F^r \rangle_t. \end{aligned}$$

Undoing the hub sign flip yields the exact identity

$$\mathbb{E}\left[f(G_1) f(G_2) \prod_{j \text{ selected leaf}} f(G_j) \right] = -\langle F^2, F^r \rangle_t. \tag{56}$$

This derivation applies directly to every bounded measurable odd f. It uses bounded conditional expectations and the absolutely convergent Mehler expansion. All the Fock products are finite because $\|F^j\|_t \le \|F\|_{jt}^j \le 1$ for $jt \le 1$. Consequently

$$L_0(f) = 1 - \sum_{h=1}^{n/2} \binom{n}{2h} \langle F^2, F^{2h} \rangle_t. \tag{57}$$

Theorem 12.2 (Two-hub negative stability)**.** *For every finite N and every measurable odd $f : \mathbb{R}^N \to [-1, 1]$,*

$$L_0(f) \le 1 - c\|F^2\|_t^2 \le 1 - cS_t(f)^2. \tag{58}$$

The same statement holds with any $0 < t' \le t$ in place of t, keeping $c = n(n-1)/4$ unchanged.

Proof. The $h = 1$ term of (57) is $-\binom{n}{2}\|F^2\|_t^2 = -2c\|F^2\|_t^2$. By Cauchy–Schwarz and (53), the absolute contribution of all remaining terms is at most $T\|F^2\|_t^2$, where

$$T = \sum_{h=2}^{n/2} \binom{n}{2h} 2^{h-1} h^h t^{h-1}.$$

Since $(2h)! \ge h! h^h$, and with $z = 2n^2 t = 1/8$,

$$T \le n^2 \sum_{h \ge 2} \frac{z^{h-1}}{h!} \le \frac{n^2 z e^z}{2} \le \frac{n^2}{14} \le \frac{n(n-1)}{4} = c.$$

Here $h! \ge 2(h-2)!$ gives the second inequality, $e^{1/8} \le (1 - 1/8)^{-1} = 8/7$ gives the third, and $4n \le 14(n-1)$ for $n \ge 2$ gives the last. For $n = 2$, the finite tail sum is empty. The tail therefore offsets at most half of the leading negative term, proving the first inequality in (58); equation (54) proves the second. Reducing t only decreases T and preserves every norm-finiteness condition, proving the last assertion. □

12.3 A global test mixture and regularization

For the positive covariance (49), write $G_i = \sqrt{t} Z + \sqrt{1-t} Z_i$ with independent standard Gaussian vectors. Let $g = T_{\sqrt{t}} f$ be the Ornstein–Uhlenbeck smoothing of f. Then g is odd, $|g| \le 1$, and $\mathbb{E} g^2 = S_t(f)$. Thus

$$L_+(f) := \mathbb{E} \prod_i (1 + f(G_i)) = \mathbb{E}(1+g)^k \le 1 + \left(2^{k-1} - 1\right) S_t(f) = 1 + C S_t(f). \tag{59}$$

Indeed the odd moments vanish, and every even moment of order at least two is at most $\mathbb{E} g^2$. Select the entire two-hub test with probability $1 - a$ and the entire positive-correlation test with probability a. With $u = S_t(f)$, completing a square gives

$$\begin{aligned} (1-a) L_0(f) + a L_+(f) &\le 1 - (1-a) c u^2 + a C u \\ &\le 1 + \frac{a^2 C^2}{4(1-a)c} = 1 + \frac{x}{4(1-a)} \le 1 + \frac{x}{2}. \end{aligned} \tag{60}$$

We used $aC^2 = cd$, $x = ad$, and (43). This is uniform over all finite dimensions and all bounded measurable odd functions. The same function must be used in both branches. The bound also holds

when the common function is chosen using auxiliary randomness independent of the branch and the queries: condition on that randomness and apply the same inequality.

To give the Boolean laws full support, let U be the uniform law on $\{-1,1\}^k$ and set

$$\mu_b = (1-\xi)D_b + \xi U \quad (b \in \{0,+\}), \qquad t' = (1-\xi)t. \tag{61}$$

Every atom has mass at least $\xi 2^{-k} > 0$, all one-bit marginals stay uniform, and the off-diagonal covariances become $t'A_*$ and $t'(J-I)$, respectively. Theorem 12.2 holds at t', and (59) holds with exactly the same $S_{t'}(f)$. Therefore (60) remains valid for the regularized covariance matrices. Define

$$\begin{aligned} \beta &= 2^{-k}\left(1 + \frac{x}{4(1-a)}\right), \\ \gamma &= (1-\xi)\big((1-a)q_0 + aq_+\big) + \xi 2^{-k} = \frac{(1-\xi)(1+x)}{k+2} + \xi 2^{-k}. \end{aligned} \tag{62}$$

These are rational. The Gaussian mixture has acceptance at most β. A common dictator has Boolean acceptance exactly γ. As $\xi(1+x) = x/8$,

$$\gamma \geq \frac{1 + 7x/8}{k+2}, \qquad \beta \leq 2^{-k}(1 + x/2). \tag{63}$$

Regularization is carried out inside each branch. The branch is selected only once for the entire N-coordinate test, followed by conditionally independent columns. This sampling order preserves the whole-test mixture in (60).

13 A common-function Boolean to Gaussian transfer

For a cube function $h : \{-1,1\}^R \to \mathbb{R}$, write its Walsh expansion as $h(z) = \sum_{S \subseteq [R]} \widehat{h}(S) \prod_{i \in S} z_i$ and define

$$T_\theta h = \sum_S \theta^{|S|} \widehat{h}(S) \prod_{i \in S} z_i, \qquad \mathrm{Inf}_i(h) = \sum_{S \ni i} \widehat{h}(S)^2.$$

A function $h : \{-1,1\}^R \to [0,1]$ is folded when $h(-z) = 1 - h(z)$. The same notation T_θ on Gaussian space denotes the Ornstein–Uhlenbeck operator, which multiplies Hermite degree j by θ^j.

Lemma 13.1 (Uniform common-function transfer)**.** *Let μ_b be finitely many fixed laws on $\{-1,1\}^k$ with uniform one-bit marginals and every atom at least $\alpha_* > 0$. Let p_b be fixed branch weights and $\Sigma_b = \mathbb{E}_{\mu_b} XX^{\mathsf{T}}$. Suppose that for every finite dimension and every common measurable odd f with $|f| \leq 1$, the corresponding Gaussian mixture satisfies*

$$\sum_b p_b \, 2^{-k} \mathbb{E}_{\Sigma_b} \prod_{j=1}^{k} (1 + f(G_j)) \leq \beta.$$

For every $\varepsilon > 0$ there are $\eta, \tau > 0$, independent of R, such that every folded $h : \{-1,1\}^R \to [0,1]$ with

$$\max_i \mathrm{Inf}_i(T_{1-\eta} h) \leq \tau \tag{64}$$

has Boolean global-branch acceptance at most $\beta + \varepsilon$. Here one first chooses b, then samples R independent columns from μ_b, obtaining row strings $X_1, \dots, X_k$, and the acceptance is $\mathbb{E} \prod_j h(X_j)$.

The proof applies correlated-space smoothing and invariance methods [Mos10, DM13] to the same odd clipped polynomial in every covariance branch. We may assume $0 < \varepsilon < 1$.

13.1 Full support gives dimension-independent smoothing

Fix a branch and a row bit X in one column, and let Y be the other $k-1$ bits. Each of the two conditional probabilities $\Pr(X = 1 \mid Y = y)$ and $\Pr(X = -1 \mid Y = y)$ is at least α_*: its numerator is an atom of mass at least α_* and its denominator is at most one. Since X is unbiased, its mean-zero function space is one-dimensional. Hence its maximal correlation with Y is

$$\|\mathbb{E}[X \mid Y]\|_2 \le 1 - 2\alpha_* =: r < 1.$$

For R independent columns, conditional expectation from the chosen row to all other rows sends a Walsh monomial of degree m to a product of independent conditional means with norm at most r^m. Images of distinct monomials are orthogonal: an index in their symmetric difference contributes one mean-zero conditional mean, independently of the other columns. Thus, for $\theta = 1 - \eta$,

$$\|\mathbb{E}\left[(h - T_\theta h)(X) \mid \text{other rows}\right]\|_2 \le \sup_{m \ge 1} (1 - \theta^m) r^m \, \|h - \mathbb{E}h\|_2. \tag{65}$$

Choose $\delta = \varepsilon/(4k)$, an integer $L \ge 1$ with $r^L \le \delta$, and $0 < \eta \le \delta/L$. For $m \le L$, $1 - \theta^m \le m\eta \le \delta$; for $m > L$, $r^m \le \delta$. Since $h \in [0, 1]$, the right side of (65) is at most δ. The product of the other row functions takes values in $[0, 1]$. Cauchy–Schwarz therefore bounds the change in acceptance from smoothing one row by δ. Smooth rows successively. In every branch the total loss is at most $k\delta = \varepsilon/4$. The same η works in all branches, and noise preserves both folding and the range $[0, 1]$.

13.2 One odd polynomial and its clipping

Set $g = 2h - 1$ and $q = T_\theta g$. These functions are odd and bounded by one. Let q_D be the truncation of q to Walsh terms of degree at most D. It is an odd multilinear polynomial, both on Boolean and Gaussian space, with exactly the same coefficients in every branch and row. Orthogonality gives

$$\begin{aligned} &\|q - q_D\|_2 \le \theta^{D+1}, && \|q_D\|_2 \le 1, \\ &\operatorname{Inf}_i(q_D) \le \operatorname{Inf}_i(q) = 4 \operatorname{Inf}_i(T_\theta h) \le 4\tau, && \sum_i \operatorname{Inf}_i(q_D) \le D. \end{aligned} \tag{66}$$

Let $c_0(z) = \max(-1, \min(1, z))$. Clipping is odd and nonexpansive, and q is $[-1, 1]$-valued on the cube, so

$$\|q - c_0(q_D)\|_2 \le \theta^{D+1}.$$

Write $P(z_1, \dots, z_k) = 2^{-k} \prod_j (1 + z_j)$. Each partial derivative of P has magnitude at most $1/2$ on $[-1, 1]^k$. The telescoping product bound and the uniform row marginals show that replacing q by $c_0(q_D)$ in the Boolean acceptance costs at most $k\theta^{D+1}/2$, regardless of within-column correlations. Choose D so that this is at most $\varepsilon/4$.

13.3 A vector Lindeberg comparison with explicit error

Choose a smooth even nonnegative mollifier of integral one supported on $[-1, 1]$, and convolve c_0 with its λ-scale version to obtain c_λ. Then c_λ is odd, takes values in $[-1, 1]$, and $\|c_0 - c_\lambda\|_\infty \le \lambda$. Define

$$\Psi(z_1, \dots, z_k) = P(c_\lambda(z_1), \dots, c_\lambda(z_k)).$$

All third partial derivatives are bounded in absolute value by $B = B(k, \lambda) < \infty$; for $0 < \lambda \le 1$ one may take $B = O_k(\lambda^{-2})$, because c_0 is one-Lipschitz and smoothing bounds its first three derivatives

by constants times $1, \lambda^{-1}, \lambda^{-2}$. Replacing c_0 by c_λ in a Boolean or Gaussian expectation costs at most $k\lambda/2$.

Fix a branch. Replace its R independent Boolean column vectors one at a time by independent Gaussian column vectors with covariance Σ_b. At coordinate i, the vector of row-polynomial evaluations has the form

$$\big(q_D(\text{row } 1), \dots, q_D(\text{row } k)\big) = A + (Z_1 B_1, \dots, Z_k B_k),$$

where A and the B_j depend only on the other columns. The replaced column is independent of them. Its old and new first moments and full covariance matrices agree. Thus, conditional on the other columns, the constant, linear, and quadratic terms in the Taylor expansion of Ψ at A cancel exactly.

Each B_j has degree at most $D-1$. Its individual row inputs are independent unbiased bits and standard Gaussians. The standard 2-to-3 hypercontractive inequality, valid also for such mixtures [Mos10, Propositions 3.14–3.15 and Theorem 3.16], gives, with $I_i = \mathrm{Inf}_i(q_D)$,

$$\mathbb{E}|B_j|^3 \le 2^{3(D-1)/2} I_i^{3/2}.$$

The same I_i applies to every row; orthogonality of distinct multilinear monomials in each mixed row gives $\mathbb{E}B_j^2 = I_i$. This bound holds even when the different B_j are dependent. The absolute third moment of a standard Gaussian is less than 2, and that of an unbiased bit is 1. Taylor's remainder is at most $(B/6)(\sum_j |Z_j B_j|)^3$. Using $(\sum_j |u_j|)^3 \le k^2 \sum_j |u_j|^3$ on both old and new columns, the ith replacement error is at most

$$\frac{Bk^3}{2}\, 2^{3(D-1)/2} I_i^{3/2}.$$

The sum of these errors is controlled by the total influence: $\sum_i I_i^{3/2} \le (\max_i I_i)^{1/2} \sum_i I_i \le 2D\sqrt{\tau}$, by (66). This removes the dependence on the number R of column replacements. The full replacement error is therefore at most

$$Bk^3\, 2^{3(D-1)/2} D\sqrt{\tau}. \tag{67}$$

The bound is independent of R, uniform over branches, and valid also for singular Σ_b. Set $\lambda = \varepsilon/(4k)$ and choose $\tau > 0$ small enough that (67) is at most $\varepsilon/4$. The two errors from smoothing the clipping function together are at most $\varepsilon/4$. Therefore the Boolean acceptance of $c_0(q_D)$ is at most its Gaussian acceptance plus $\varepsilon/2$.

13.4 Applying the Gaussian bound to one common function

Take $H(z) = c_0(q_D(z))$. This is one measurable odd $[-1,1]$-valued function in dimension R, independent of the branch. The lemma's Gaussian premise applies to it. The errors are $\varepsilon/4$ from row smoothing, $\varepsilon/4$ from Boolean truncation and clipping, and $\varepsilon/2$ from the comparison just proved. The original acceptance is therefore at most $\beta + \varepsilon$, proving Lemma 13.1. The parameters are chosen in the following order:

$$\varepsilon \;\longrightarrow\; \delta, L, \eta \;\longrightarrow\; D \;\longrightarrow\; \lambda, B \;\longrightarrow\; \tau.$$

All parameters are independent of R, and the influence hypothesis applies to the smoothed function.

14 The Unique Games reduction

We now apply Lemma 13.1 to the two full-support laws (61), with branch weights $1-a, a$ and Gaussian bound (62). We use the standard loop-free, undirected, regular form of the Unique Games

gap, as in the folded reduction of De and Mossel [DM13, Section 5]. The alphabet is $[R]$. Degree counts edge multiplicities, and a uniform neighbor is sampled by a uniform incident-edge occurrence. If edges are parallel, the chosen edge occurrence is part of the notation (v, w) and determines its permutation. Thus ordered neighbor tuples below mean ordered incident-edge tuples. Write $\pi_{v,w} : [R] \to [R]$ for the edge permutation taking labels at w to labels at v, so an edge is satisfied when

$$L(v) = \pi_{v,w}(L(w)), \qquad \pi_{w,v} = \pi_{v,w}^{-1}.$$

The Unique Games Conjecture supplies a fixed alphabet for any fixed positive completeness loss and soundness bound. These parameters will be chosen after the dimension-free transfer parameters.

14.1 Folding and the verifier

For each vertex w, the proof is a Boolean table $h_w : \{-1, 1\}^R \to \{0, 1\}$ with folding $h_w(-z) = 1 - h_w(z)$. Implement folding by storing one Boolean variable for each antipodal pair, using the representative with first coordinate $+1$. A query at the opposite representative uses the negated literal. Thus every assignment to the stored CSP variables defines a folded table. Each CSP variable is identified by its proof vertex (or clone) and its antipodal query pair. Both branches use the same graph, tables, and variable identifiers; the branch choice changes only the query law. For a center v and its neighbor w, define

$$F_{v,w}(z) = h_w(y), \qquad y_j = z_{\pi_{v,w}(j)}. \tag{68}$$

If h_w is the dictator at label $L(w)$ and the edge is satisfied, then $F_{v,w}$ is the dictator at $L(v)$. Moreover,

$$\mathrm{Inf}_i(T_\theta F_{v,w}) = \mathrm{Inf}_{\pi_{v,w}^{-1}(i)}(T_\theta h_w). \tag{69}$$

The verifier samples a uniform center v, then k neighbors $w_1, \ldots, w_k$ independently and with replacement. It selects one global branch b, samples R conditionally independent columns from μ_b, and accepts the conjunction of the k literals $F_{v,w_j}(X_j)$. This directly produces signed Max-k-AND clauses.

14.2 Completeness

Suppose a labeling satisfies at least $1 - \delta_{\mathrm{UG}}$ of the edges, and use its dictator tables. A union bound shows that all k selected edges are satisfied with probability at least $1 - k\delta_{\mathrm{UG}}$. Conditional on this event all pulled-back tables are the same dictator, so branch-averaged acceptance is exactly γ from (62). Thus the verifier accepts with probability at least

$$(1 - k\delta_{\mathrm{UG}})\gamma \geq \gamma - k\delta_{\mathrm{UG}}. \tag{70}$$

This gives completeness close to $1/(k + 2)$.

14.3 Neighborhood averaging gives a common function

For an arbitrary assignment, define the folded $[0, 1]$-valued function

$$g_v(z) = \mathbb{E}_{w \text{ uniform neighbor of } v} F_{v,w}(z).$$

Conditionally on the center, the branch, and all query strings, independent neighbor sampling gives the exact identity

$$\mathbb{E}_{w_1,\ldots,w_k} \prod_{j=1}^{k} F_{v,w_j}(X_j) = \prod_{j=1}^{k} g_v(X_j). \tag{71}$$

This identity expresses both branches in terms of the same g_v, as required by the transfer lemma.

14.4 Influence decoding

Fix $0 < \varepsilon < 1$ and use Lemma 13.1 with error $\varepsilon/2$, obtaining $\eta, \tau > 0$ and $\theta = 1 - \eta$. We may decrease τ to ensure $\tau \le 1$. If total verifier acceptance exceeds $\beta + \varepsilon$, a fraction at least $\varepsilon/2$ of centers have acceptance greater than $\beta + \varepsilon/2$. Indeed, if their fraction is p, the total is at most $\beta + \varepsilon/2 + p$, since individual acceptances lie in $[0, 1]$. Each such center v has a coordinate i with

$$\mathrm{Inf}_i(T_\theta g_v) > \tau.$$

Noise commutes with permutations. Jensen's inequality on Fourier coefficients, whose squared noise multiplier is $\theta^{2|S|}$, and (69) imply

$$\mathrm{Inf}_i(T_\theta g_v) \le \mathbb{E}_w \, \mathrm{Inf}_{\pi_{v,w}^{-1}(i)}(T_\theta h_w). \tag{72}$$

Since each influence is at most one, at least a $\tau/2$ fraction of neighbors have the corresponding influence at least $\tau/2$. For if that fraction is p, the average is at most $p + \tau/2$, forcing $p > \tau/2$.

At every vertex u define

$$A(u) = \{i : \mathrm{Inf}_i(T_\theta h_u) \ge \tau/2\} \cup \{i : \mathrm{Inf}_i(T_\theta g_u) \ge \tau\}.$$

The first set records candidate labels indicated by the table at u; the second records candidates indicated by its neighbors after their labels are transported to u. Including both sets places the influential coordinate of a good center and its transported coordinate at a good neighbor in their respective candidate sets. For any $[0, 1]$-valued function f,

$$\sum_i \mathrm{Inf}_i(T_\theta f) = \sum_S |S| \theta^{2|S|} \widehat{f}(S)^2 \le \frac{1}{\eta} \|f\|_2^2 \le \frac{1}{\eta}.$$

The multiplier bound follows from $m\theta^{2m} \le m\theta^m \le \sum_{j=0}^{m-1} \theta^j \le 1/(1-\theta) = 1/\eta$. Hence

$$|A(u)| \le \frac{3}{\eta\tau} \le \frac{4}{\eta\tau}.$$

Choose independent uniform labels from each nonempty $A(u)$ and arbitrary labels otherwise. For each of the good directed edges identified above, the coordinate i lies in $A(v)$ and $\pi_{v,w}^{-1}(i)$ lies in $A(w)$. Since the graph is loop-free, the two choices are independent, and this edge is satisfied with probability at least $(\eta\tau/4)^2$. Averaging over the good centers and good neighbors gives expected Unique Games value at least

$$\kappa = \frac{\varepsilon \eta^2 \tau^3}{64} > 0. \tag{73}$$

Here $\kappa = (\varepsilon/2)(\tau/2)(\eta\tau/4)^2$: the factors account for the fraction of good centers, the fraction of good neighbors, and the two independent label choices. Some deterministic labeling therefore attains at least this value. The constant is independent of R. Choose the Unique Games soundness strictly below κ and its completeness loss small enough for (70). The alphabet R is then fixed by UGC. The transfer and decoding parameters have all been fixed before the alphabet is chosen.

14.5 Distinct variables and exact rational multiplicities

If constraints must use k distinct variables, first replace every original Unique Games vertex by M clones, and every edge by all M^2 corresponding clone edges with the same permutation. The graph remains regular and loop-free. Completeness is preserved. Soundness is preserved as well: for any

labeling of the clone graph, choose one independent uniform clone label at each original vertex. The expected original edge value equals the clone edge value, so some original labeling achieves at least that value. In the verifier on the clone graph the probability that two selected neighbor vertices coincide is at most $\binom{k}{2}/M$. To see this, each fixed clone neighbor has probability at most $1/M$, even when the original graph has parallel edges; for any pair of independently sampled neighbors the collision probability is therefore at most $1/M$. Apply (71) to the unfiltered verifier, then remove every clause with a repeated proof vertex. Distinct proof vertices use disjoint long-code variable sets, so the retained clauses have k distinct variables regardless of the query strings and folding signs.

Let p be the total mass removed. If $p \le 1/2$, conditioning on retention changes the acceptance of any fixed assignment by at most $p/(1-p) \le 2p$. Indeed, the unconditioned value is $(1-p)v_{\mathrm{keep}} + pv_{\mathrm{remove}}$ with both conditional values in $[0,1]$. The same uniform estimate applies to the optimum. Taking the constant M sufficiently large absorbs this loss into any prescribed final gap slack. Conditioning on distinct neighbors makes the neighbor choices dependent. The uniform loss estimate transfers the already established completeness and soundness bounds to these conditioned tuples.

All slice probabilities, densities, regularization weights, and branch weights are rational. For fixed k and final accuracy, both R and M are constants. On a regular graph with V vertices and degree d_{UG}, enumerate every center, every ordered incident-edge k-tuple, both branches, and every R-column matrix. There are at most

$$2Vd_{\mathrm{UG}}^{k}2^{kR}$$

such choices. Let D_0^{den} clear the denominators of both one-column laws and let A^{den} clear the denominators of the branch weights. Multiplication of every verifier probability by

$$Vd_{\mathrm{UG}}^{k}A^{\mathrm{den}}(D_0^{\mathrm{den}})^{R}$$

gives integer clause multiplicities. Their total equals this quantity before deletion, because the probabilities sum to one. It is polynomial in the Unique Games input size, with constants depending on k and the requested gap. Deleting collisions removes integer copies and renormalizes the uniform retained instance. All multiplicities remain polynomial in the input size for each fixed arity and gap. This yields an unweighted instance with exactly k distinct variables per clause.

14.6 Conclusion of the hardness proof

Given any sufficiently small final gap slack $\zeta > 0$, use soundness error $\varepsilon < \zeta/2$ in the verifier analysis. Choose M so that the deletion loss is less than $\zeta/2$, and choose the Unique Games completeness loss so that $k\delta_{\mathrm{UG}} < \zeta/2$. The transfer parameters depend only on the fixed soundness error, and UGC then supplies a fixed alphabet with soundness below (73). Completeness is at least $\gamma - \zeta$ and soundness at most $\beta + \zeta$. Every approximation factor strictly above β/γ distinguishes this gap for sufficiently small ζ. Finally (63) gives

$$\frac{\beta}{\gamma} \le \frac{k+2}{2^k}\frac{1+x/2}{1+7x/8} = R_k < \frac{k+2}{2^k},$$

proving Theorem 1.4. Since Max-k-AND is a special case of general Boolean Max-k-CSP, the stated implication for the latter follows.

15 Discussion

The common reference vector makes a k-literal conjunction a two-orthant event in dimension $k+1$. Mulgund's and Bakshi's Gaussian inequalities, together with the scalar estimate proved here, yield

the factor α_k for $k \geq 5$. The all-sign-pattern instance shows that this is the exact uniform factor of the relaxation. Combined with known hardness, it settles the deterministic polynomial-time threshold for every odd $k \geq 5$ under UGC and P $\neq$ NP, and for $k = 2^r - 1 \geq 7$ under P $\neq$ NP.

The even-arity construction uses a distinct Gaussian structure. Oppositely correlated hubs produce a negative leading mixed moment, and the Fock product estimate controls the higher moments relative to it, in every finite dimension. Mixing whole tests gives a linear completeness gain while increasing soundness only quadratically. The common-function transfer and neighborhood averaging extend this estimate to arbitrary folded low-influence decisions.

For even $k \geq 6$, the remaining interval is $[\alpha_k, R_k]$, with arbitrarily small slack. Determining the exact ratio in this interval remains open. The case $k = 4$ also retains an algorithmic question alongside the new hardness bound. On the implementation side, deterministic rounding has polynomial dependence on $1/\varepsilon$, while randomized rounding has polynomial dependence on $\log(1/\varepsilon)$.

A An exact arithmetic certificate

The endpoint proof reduces the estimate involving Φ and the parameter q to explicit rational inequalities, each checkable by integer arithmetic. The accompanying script `certify_endpoint.py` uses Python's exact `Fraction` type. It verifies the elementary arithmetic in Section 5. The analytic argument establishes the reductions and Taylor remainder signs and uses the classical bounds $25/8 < \pi < 22/7$.

In particular, it checks the squared bounds on $(2\pi)^{-1/2}$, the alternating degree-three and degree-four bounds for $e^{-1/8}$, the integral lower bound for $\Phi(1/2)$, the positive logarithmic series lower sum, and both strict margins

$$G_6 - \frac{3}{10} > \frac{1367524139}{1604608399800} > 0, \qquad e^{9/5} - 6 > \frac{42629}{1250000} > 0.$$

The first margin follows directly from the strict rational lower bound in (27). Running the script prints the following exact output:

```
All rational checks passed.
G_6 - 3/10 > 1367524139/1604608399800
exp(9/5) - 6 > 42629/1250000
```

AI methodology

OpenAI Codex assisted with the abstract and introduction, English exposition throughout the paper, elaborating and checking proof details, LaTeX preparation, reference verification, and consistency checks. The authors retain responsibility for the mathematical claims and the final manuscript.

References

[AH12] Per Austrin and Johan Håstad. *On the Usefulness of Predicates*. arXiv:1204.5662v1, 2012. arxiv.org/abs/1204.5662v1.

[AK13] Per Austrin and Subhash Khot. *A Characterization of Approximation Resistance for Even k-Partite CSPs*. arXiv:1301.2731, 2013. arxiv.org/abs/1301.2731.

[AM09] Per Austrin and Elchanan Mossel. Approximation resistant predicates from pairwise independence. *Computational Complexity*, 18(2):249–271, 2009. doi:10.1007/s00037-009-0272-6.

[Bak26] Ainesh Bakshi. *On the Approximability of Boolean Max-k-CSP*. arXiv:2608.05331v1, 2026. arxiv.org/abs/2608.05331v1.

[Cha12] Siu On Chan. *Approximation Resistance from Pairwise Independent Subgroups*. Electronic Colloquium on Computational Complexity, Report TR12-110, 2012. ECCC TR12-110.

[CMM09] Moses Charikar, Konstantin Makarychev, and Yury Makarychev. Near-optimal algorithms for maximum constraint satisfaction problems. *ACM Transactions on Algorithms*, 5(3), Article 32, pages 1–14, 2009. doi:10.1145/1541885.1541893. Author-hosted full text.

[DPLS11] Paolo Da Pelo, Alberto Lanconelli, and Aurel I. Stan. *A Hölder–Young–Lieb Inequality for Norms of Gaussian Wick Products*. arXiv:1101.2947, 2011. arxiv.org/abs/1101.2947.

[dKV16] Etienne de Klerk and Frank Vallentin. On the Turing model complexity of interior point methods for semidefinite programming. *SIAM Journal on Optimization*, 26(3):1944–1961, 2016. doi:10.1137/15M103114X. Preprint: arXiv:1507.03549v2.

[DM13] Anindya De and Elchanan Mossel. *Explicit Optimal Hardness via Gaussian Stability Results*. arXiv:1202.5258v6, 2013. arxiv.org/abs/1202.5258v6.

[GLS88] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Algorithms and Combinatorics, vol. 2. Springer, 1988. doi:10.1007/978-3-642-97881-4.

[GW95] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995. doi:10.1145/227683.227684.

[Has05] Gustav Hast. Approximating Max kCSP—outperforming a random assignment with almost a linear factor. In *Automata, Languages and Programming (ICALP 2005)*, Lecture Notes in Computer Science, vol. 3580, pages 956–968. Springer, 2005. doi:10.1007/11523468_77.

[KTW13] Subhash Khot, Madhur Tulsiani, and Pratik Worah. *A Characterization of Approximation Resistance*. arXiv:1305.5500v2, 2013. arxiv.org/abs/1305.5500v2.

[KSS02] Hui-Hsiung Kuo, Kimiaki Saitô, and Aurel I. Stan. A Hausdorff–Young inequality for white noise analysis. In T. Hida and K. Saitô, editors, *Quantum Information IV*, pages 115–126. World Scientific, 2002. doi:10.1142/9789812777324_0007. The L^2 product estimate is reproduced and attributed in [DPLS11].

[MR99] Sanjeev Mahajan and H. Ramesh. Derandomizing approximation algorithms based on semidefinite programming. *SIAM Journal on Computing*, 28(5):1641–1663, 1999. doi:10.1137/S0097539796309326.

[MM14] Konstantin Makarychev and Yury Makarychev. Approximation algorithm for non-Boolean Max-k-CSP. *Theory of Computing*, 10(13):341–358, 2014. doi:10.4086/toc.2014.v010a013.

[Mak26] Yury Makarychev. *Sharp Analysis of Gaussian Rounding for Boolean Max k-CSP*. arXiv:2608.07800v1, 2026. arxiv.org/abs/2608.07800v1.

[Mos10] Elchanan Mossel. Gaussian bounds for noise correlation of functions. *Geometric and Functional Analysis*, 19(6):1713–1756, 2010. doi:10.1007/s00039-010-0047-x. Preprint: arXiv:math/0703683v6.

[Mul26] Abhijeet Mulgund. *Stochastic Domination of Gaussian Maxima: A Resolution of the Weak Simplex Conjecture*. arXiv:2607.14087v4, 2026. arxiv.org/abs/2607.14087v4.

[Rag08] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC 2008)*, pages 245–254. ACM, 2008. doi:10.1145/1374376.1374414.

[Tre98] Luca Trevisan. Parallel approximation algorithms by positive linear programming. *Algorithmica*, 21(1):72–88, 1998. doi:10.1007/PL00009209.