Polynomial compatibility on the Boolean cube from disjoint-support separators

Alec Kriebel
ORCID: 0009-0001-9320-500X

17 October 2026

Abstract

We give an explicit affirmative answer to the two-polynomial question posed by Austrin in Oberwolfach Report 9/2025. Two nonzero real multilinear polynomials of degree at most d on the uniform Boolean cube have a common nonzero vertex whenever every relative coordinate influence is at most $1/(4096d^8)$. In fact, disjoint supports force a relative influence at least $1/(1024d^8)$. The result is a deduction from the disjoint-support separator of Li–Li–Li–Liu and the covariance inequality of O'Donnell–Saks–Schramm–Servedio, combined with a square-density estimate. We supply the complete deduction and an explicit separator constant. No novelty or first-priority claim is made for this consequence.

1 The question and the quantitative answer

All expectations refer to independent uniform signs $X \in \{-1,1\}^n$. For a nonzero real multilinear polynomial $a(x) = \sum_{S\subseteq[n]} \widehat{a}(S) \prod_{i\in S} x_i$, set

$$\|a\|_2^2 = \mathbb{E}a(X)^2 = \sum_S \widehat{a}(S)^2, \qquad \mathrm{RelInf}_i(a) = \frac{\sum_{S\ni i} \widehat{a}(S)^2}{\|a\|_2^2}.$$

The denominator includes the constant coefficient. Austrin's question [1, p. 445] asks whether an inverse-polynomial influence bound in the degree guarantees a common nonzero vertex for two such polynomials. The following gives that guarantee with absolute constants.

Theorem 1. *Let $n, d \geq 1$ be integers and let f, g be nonzero real multilinear polynomials on $\{-1,1\}^n$, each of degree at most d. If $f(x)g(x) = 0$ at every vertex, then*

$$\max_{i\in[n]} \max\{\mathrm{RelInf}_i(f), \mathrm{RelInf}_i(g)\} \;\geq\; \frac{1}{1024d^8}. \tag{1}$$

Consequently, if every relative influence of both polynomials is at most $1/(4096d^8)$, some vertex satisfies $f(x)g(x) \neq 0$.

The structural input is the prior separator of Li, Li, Li and Liu [3, Lemma 3.5]; the covariance input is already [5, §3.3, Eq. (6)]. Neither is a new theorem of this note. The Markov/Bernoulli presentation in [4, Fact 5 and Corollary 1] makes the numerical constant transparent. The resulting target-specific application is written out below, including arbitrary common zeros. The general distribution formulations in [2, Conjectures 1.2 and 5.5] and other groups or input measures are outside the theorem asserted here.

2 A quantitative total separator

We reconstruct the prior separator to keep both its constant and its behavior on common zeros explicit. Here alone we use $\{0,1\}$ coordinates; the affine change between the two cubes preserves degrees and query depths.

Lemma 1 (Li–Li–Li–Liu separator, with an extracted constant)**.** *If real multilinear p, q of degree at most $d \geq 1$ on $\{0,1\}^n$ satisfy $pq = 0$ everywhere, there is a total deterministic decision tree of depth at most $16d^4$, with output $h \in \{0,1\}$, such that $h = 1$ on $\operatorname{supp} p$ and $h = 0$ on $\operatorname{supp} q$. No coordinate repeats along a branch.*

Proof. We first recall the block bound used in [3, 4]. Suppose $R : \{0,1\}^b \to \mathbb{R}$ has degree at most m, satisfies $|R| \leq 1$, $R(0) = 1$, and $R(e_j) \leq 0$ for all j. Its expectation $P(t)$ at independent Bernoulli(t) bits is a polynomial of degree at most m, bounded by one throughout $[0,1]$. Moreover,

$$P'(0) = \sum_{j=1}^{b} (R(e_j) - 1) \leq -b.$$

The classical Markov inequality, rescaled to $[0,1]$, gives $|P'(0)| \leq 2m^2$. Thus $b \leq 2m^2$.

At a node, restrict p, q by the coordinates already queried. If one restriction is zero, label for the other side, with an arbitrary label if both are zero. Otherwise neither restriction can be a nonzero constant, since the other has disjoint support. Put $D = \max(\deg p, \deg q) \geq 1$. The multilinear representative of $s = p^2 - q^2$ has degree at most $2D$. Choose a vertex z maximizing $|s|$, whose maximum is positive. Suppose $p(z) \neq 0$, so $s(z) > 0$ and $q(z) = 0$.

Choose an inclusion-maximal family of pairwise disjoint supports $M_1, \ldots, M_b$ of maximum-degree monomials of q. For each M_j, fixing all variables outside it to their values at z preserves its top coefficient: a further contribution would require a strictly containing monomial. This restriction is therefore nonzero, although it vanishes at z. Some nonempty $E_j \subseteq M_j$ can be flipped from z to make $q \neq 0$. There $p = 0$ and $s < 0$. The blocks E_j are disjoint.

Encode these block flips by b Boolean variables and divide s by $|s(z)|$. The resulting function R is bounded by one, has $R(0) = 1$ and $R(e_j) < 0$, and has degree at most $2D$: each original coordinate becomes an affine function of one block bit, and multilinearization cannot increase degree. The block bound gives $b \leq 8D^2$.

The union $H = \bigcup_j M_j$ contains at most $8d^3$ coordinates and hits every maximum-degree monomial of q, by maximality. Query all coordinates in H. On every answer branch, q becomes zero or strictly loses degree; the degree of p cannot increase. The case $q(z) \neq 0$ is symmetric. Each nonterminal stage decreases $\deg p + \deg q$ by at least one, from an initial value at most $2d$. Thus at most $2d$ stages, each costing at most $8d^3$ queries, suffice. Only free coordinates are queried. Zero restrictions terminate before their degree is needed. These rules give a total tree, including on the common-zero set. □

3 Covariance and squared densities

For a real function u on the sign cube, let

$$A_i(u) = \mathbb{E}|u(X) - u(X^{(i)})|,$$

where $X^{(i)}$ independently resamples coordinate i; it leaves that coordinate unchanged with probability one-half.

Lemma 2 (OSSS covariance inequality)**.** *Let a total deterministic tree, without repeated branch queries, compute $h : \{-1,1\}^n \to \{0,1\}$. If r_i is the probability that it queries coordinate i, then for every real u,*

$$|\operatorname{Cov}(h,u)| \le \frac{1}{2}\sum_i r_i A_i(u). \tag{2}$$

Proof. This is directly [5, §3.3, Eq. (6)], applied to $2h-1$. We give the finite-cube argument for completeness. Run the tree solely on X, and take an independent uniform Y. Let Z_t replace the first t queried coordinates of X by those of Y, and let L be the branch length. A complete transcript is precisely a cylinder fixing the queried X-coordinates. Conditional on it, Z_L is uniform: queried coordinates come from Y, and unqueried X-coordinates remain independent fair signs. It is therefore independent of the transcript and its output. Consequently

$$\operatorname{Cov}(h,u) = \mathbb{E}\left[(h(X)-\tfrac{1}{2})\sum_{t=1}^{L}\big(u(Z_{t-1})-u(Z_t)\big)\right].$$

Conditional on a *pre-query* transcript whose next index is i, the previously fixed X-coordinates have already been replaced by independent Y-coordinates. All remaining X-coordinates are unrestricted, and Y_i is fresh because queries do not repeat. Thus (Z_{t-1}, Z_t) has exactly the unconditional uniform coordinate-resampling-i law under this conditioning. Using $|h-\frac{1}{2}| = \frac{1}{2}$, taking the absolute telescoping sum, and summing the probabilities of these pre-query cylinders gives (2). All sums are finite, including variable-length branches; u need not be computed by the tree. □

Lemma 3 (Square-density estimate)**.** *For any nonzero real a, with $\rho_i = \operatorname{RelInf}_i(a)$ defined from its unique multilinear expansion,*

$$A_i\left(\frac{a^2}{\|a\|_2^2}\right) \le 2\sqrt{\rho_i(1-\rho_i)} \le 2\sqrt{\rho_i}.$$

Proof. Write $a = b + x_i c$, with b, c independent of x_i. Parseval gives $\|a\|_2^2 = \mathbb{E}b^2 + \mathbb{E}c^2$ and $\rho_i = \mathbb{E}c^2/\|a\|_2^2$. A sign flip changes a^2 by $4x_i bc$, while independent resampling flips with probability one-half. Hence the left side equals $2\mathbb{E}|bc|/\|a\|_2^2$. Cauchy–Schwarz gives the asserted bound, including $\rho_i = 0, 1$. □

Proof of Theorem 1. Assume disjoint supports and take the total separator from Lemma 1, of depth $T \le 16d^4$. Set $u = f^2/\|f\|_2^2$, $v = g^2/\|g\|_2^2$, and $\mu = \mathbb{E}h$. Then $\mathbb{E}u = \mathbb{E}v = 1$, $hu = u$, and $hv = 0$, so

$$|\operatorname{Cov}(h,u)| + |\operatorname{Cov}(h,v)| = (1-\mu) + \mu = 1.$$

These identities also hold at common zeros, where both densities vanish. With δ the maximum relative influence of either polynomial, Lemmas 2–3 give

$$1 \le \sum_i r_i\Big(\sqrt{\operatorname{RelInf}_i(f)} + \sqrt{\operatorname{RelInf}_i(g)}\Big) \le 2\Big(\sum_i r_i\Big)\sqrt{\delta} \le 2T\sqrt{\delta} \le 32d^4\sqrt{\delta}.$$

This proves (1). If $\delta \le 1/(4096d^8)$, the last upper bound is at most $1/2$, a contradiction. □

A nonzero constant has full support, so constants cause no exception. The proof also covers $n < d$. The earlier exponential sufficient bound uses $\delta < 2^{-d}/d$, as in [2, §5.2]. Its non-strict endpoint printed in [1] fails already for $d = 1$, $f = 1 + x_1$, $g = 1 - x_1$, both of relative influence $1/2$. The smaller inclusive threshold in Theorem 1 has slack. No optimality of its exponent or constants is asserted.

4 Attribution, verification, and review status

The contribution of this note is the explicit application above to the stated two-polynomial question. The separator, the real covariance inequality, and the Markov/Bernoulli method are credited prior inputs, available before this note. Reconstructing their proofs does not establish novelty. A bounded primary-source audit found no earlier recorded complete target-specific deduction in the bodies inspected; it does not certify historical priority. In particular, the full ePrint version of the 2022 paper and the full 2023 non-interactive-commitment followup were inaccessible to that audit. Detailed source/version records, inspection limits, and independent priority reviews accompany the note. No firstness or current global literature-status claim is made.

The accompanying source and verification archive includes the original proof record, new independent mathematical audits, exact rational diagnostics, and a portable offline runner. Finite enumerations check adaptive trees, support geometry, Fourier identities, boundary cases, and deliberate incorrect variants in ordinary and optimized execution. They supplement the universal written proof; they are not a formal proof certificate or an exhaustive check in arbitrary dimension. Third-party source bodies are not redistributed.

AI tools were used extensively in developing, writing, reproducing and adversarially reviewing this work. The reviews are internal AI-assisted checks. This preprint is unrefereed; no external human peer review, journal acceptance, or proof-assistant certification is claimed.

References

[1] P. Austrin, *Zeros of Low-Degree, Low-Influence Polynomials on the Boolean Hypercube*, contributed problem in MFO–RIMS Tandem Workshop: Optimization, Theoretical Computer Science and Algebraic Geometry: Convexity and Beyond, Oberwolfach Reports **22** (2025), no. 1, 405–450, at p. 445. doi:10.4171/OWR/2025/9.

[2] P. Austrin, H. Chung, K.-M. Chung, S. Fu, Y.-T. Lin and M. Mahmoody, *On the Impossibility of Key Agreements from Quantum Random Oracles*, CRYPTO 2022, LNCS **13508**, 165–194. doi:10.1007/978-3-031-15979-4_6.

[3] L. Li, Q. Li, X. Li and Q. Liu, *Impossibility of Perfectly Complete Many-Round Key Agreement in the QROM*, 4 August 2026. arXiv:2608.03824v1, Lemma 3.5, pp. 7–8.

[4] R. Kothari, M. Kovacs-Deak, D. Wang and R. Z. Yang, *Rational degree is polynomially related to degree*, 28 September 2026. arXiv:2601.08727v3, Fact 5, Theorem 1 and Corollary 1.

[5] R. O'Donnell, M. Saks, O. Schramm and R. A. Servedio, *Every decision tree has an influential variable*, FOCS 2005, 31–39. doi:10.1109/SFCS.2005.34; arXiv:cs/0508071v1, §3.3, Eq. (6), p. 8.