Guarded Boolean Operators: Exact Dispatch and a Polynomial Guard-Degree Tradeoff

Brian Theory

Abstract

A finite family of continuous scores determines a Boolean operator wherever all scores are nonzero. We study exact identification of that operator under an explicit guard language, using the real polynomial $F(x, y) = x + y + xy$ on signed positive-magnitude representatives as a complete case study. Its parameter domain has exactly four operator chambers, labelled by AND, the two projections, and XNOR. In the original magnitude coordinates, let m_d be the minimum number of binary polynomial-sign tests of total degree at most d needed for exact dispatch on the zero-free domain. We prove $m_1 = \infty$, $m_2 = m_3 = 3$, and $m_d = 2$ for $d \geq 4$; moreover, both polynomials in any two-test solution must have degree at least four. Reciprocal coordinates linearize the chamber boundaries, making the coordinate dependence explicit. Three truth-table coordinate probes remain necessary in the worst case, including for adaptive identification. We place these calculations in a zero-aware certification contract, distinguish fixed test selection from per-instance and adaptive costs, and prove directly that dispatch is unchanged by faithful correspondence-matrix and fixed-frame logical-matrix representations. The general sign quotient and minimum-test-set formulation are established machinery; the contribution is a self-contained worked analysis of how coordinates, permitted observations, and zero handling change exact dispatch.

1 Introduction

A continuous model can induce different Boolean operators as its parameters vary. For example, evaluate the real score $x + y$ at $(a, b), (a, -b), (-a, b), (-a, -b)$ with $a, b > 0$, and declare positive output true and negative output false. If $a > b$, the resulting truth table is the first projection; if $b > a$, it is the second. At $a = b$, two scores vanish and the strict binary interpretation is undefined. Thus even this elementary model requires both a parameter partition and a zero policy.

The broad passage from continuous parameters to discrete logical data is not new. DSGRN constructs combinatorial parameter descriptions for switching regulatory networks and attaches dynamical information to them [12, 13, 15]. Logical bifurcation diagrams relate parameter variation in piecewise differential models to logical valuations and attractors [1]. Joint realizability and parameter adjacency of monotone Boolean functions are treated in [11], and recent work embeds monotone Boolean models into multilevel and DSGRN models [14]. These structured dynamical theories have substantially more application-specific content than a generic score-sign partition.

Likewise, strict polynomial sign conditions and their connected components are standard real algebraic geometry [4, 5, 16]; collecting sign semantics belongs to abstract interpretation [9]; and selecting a minimum distinguishing test family is a classical optimization problem [17, 6]. Predicate minimization in verification is related but can optimize preservation of a property rather than exact recovery of every operator label [8]. We use these connections to specify the problem, not

to claim a new abstraction theory, a new hardness theorem, or a replacement for existing solvers. The minimum number of polynomial inequalities describing a set is also an established question [3]. Here the optimization instead requires joint decoding of four labels on a promised domain, with a fixed degree bound and fixed original coordinates; the general polynomial-description literature is therefore relevant context, not a priority claim for the concrete calculation.

Question and main result. The question is narrower: *which observations must be available, and what must be proved about them, before a parameter value can be assigned its exact Boolean operator?* For

$$F(x,y) = x + y + xy, \qquad (x,y) = (\varepsilon_1 a, \varepsilon_2 b), \qquad a, b > 0,$$

Sections 2–3 give a complete answer for three explicit contracts. Three truth-table coordinates distinguish the four realized operators. No finite library of affine tests in (a, b) suffices. Among polynomial-sign libraries, three tests are optimal at degrees two and three, whereas two quartic tests are optimal once degree four is allowed. The last statement concerns *two polynomial observations*, not two queries to the original truth-table-coordinate oracle.

Scope. This is a mathematical methods note built around an exact example and its certification requirements. The degree tradeoff is proved by an elementary boundary-divisibility argument. The synthesis of these facts is not offered as theorem-level priority clearance for every constituent observation. In particular, the curved boundaries become straight lines under reciprocal coordinates, and fewer sign tests do not by themselves imply less arithmetic or faster execution. The accompanying implementation is an exact rational reference dispatcher and regression suite, not a general certified guard-synthesis system or an empirical performance study.

Correspondence matrices (CMs) and logical matrices (LMs) provide optional output representations. We follow the companion manuscript's terminology [19], but prove every representation fact needed here. Readers interested only in dispatch may regard a CM simply as a reshaped truth table. Classical trajectory and Walsh calibrations are collected in the appendix rather than used as novelty claims.

2 Score families and the nonlinear four-operator model

2.1 Strict semantics and operator fibers

Let Λ be a nonempty open subset of $\mathbb{R}^p$, let $n \geq 1$, and suppose

$$g_\alpha : \Lambda \longrightarrow \mathbb{R}, \qquad \alpha \in \{0,1\}^n,$$

are continuous. They may arise by evaluating a continuous model at continuously varying representatives. For signed magnitudes, the representative for α is $\varepsilon(\alpha) \odot r$, where $\varepsilon_i(\alpha) = 2\alpha_i - 1$ and $r_i > 0$.

Define the partial threshold $q(t) = 1$ for $t > 0$ and $q(t) = 0$ for $t < 0$; $q(0)$ is *undefined*. Set

$$\Delta = \bigcup_{\alpha \in \{0,1\}^n} g_\alpha^{-1}(0), \qquad \Lambda^\circ = \Lambda \setminus \Delta, \qquad \Theta_\lambda(\alpha) = q(g_\alpha(\lambda)) \quad (\lambda \in \Lambda^\circ).$$

Thus Δ is the union of all score-zero sets, not their intersection. It is closed relative to Λ. Neither nonemptiness nor density of Λ° is assumed. The numeric correspondence tensor is $[\Theta][\alpha] = \Theta(\alpha)$, in true-first order on each axis; for $n = 2$ its entries are ordered $(11, 10, 01, 00)$.

Proposition 2.1 (Elementary exact quotient)**.** *The nonempty fibers*

$$\Omega_\Theta = \{\lambda \in \Lambda^\circ : \Theta_\lambda = \Theta\} = \bigcap_{\Theta(\alpha)=1} \{g_\alpha > 0\} \cap \bigcap_{\Theta(\alpha)=0} \{g_\alpha < 0\}$$

form the unique coarsest partition of Λ° into sets carrying constant operators. Each fiber is open in Λ and clopen in Λ°. Its connected components are open and are precisely connected components of Λ°.

Proof. The displayed finite intersection is open by continuity. Its relative complement is the union of the other fibers and is open as well. Every set carrying a constant operator lies in the corresponding fiber, proving the partition claim. An open Euclidean set is locally path-connected, so its connected components are open. A connected set cannot meet two distinct clopen fibers, proving the component assertion. □

We call those components *operator chambers*. There are finitely many operator fibers but, for merely continuous scores, possibly infinitely many chambers. For instance, $g_1(t) = \sin t$ and $g_0(t) = 1$ on $\mathbb{R}$ give two fibers and infinitely many interval chambers. If Λ is open semialgebraic and the scores are polynomial, the fibers are strict sign-condition realizations and have finitely many semialgebraic connected components [4, 5, 16]. This is an application of established sign-condition theory.

2.2 Exact classification and a convex normal form

All arithmetic in the following model is ordinary real arithmetic, not arithmetic over a Boolean field. For $F(x, y) = x + y + xy$ on signed magnitudes, the four scores are

$$\begin{aligned} g_{11} &= a + b + ab, \qquad h_1 := g_{10} = a - b - ab, \\ h_2 := g_{01} &= -a + b - ab, \quad h_3 := g_{00} = -a - b + ab. \end{aligned} \tag{1}$$

Here $g_{11} > 0$ everywhere on $\Lambda = \mathbb{R}^2_{>0}$, while

$$h_1 + h_2 = -2ab < 0, \qquad h_1 + h_3 = -2b < 0, \qquad h_2 + h_3 = -2a < 0. \tag{2}$$

At most one of the h_i can therefore be positive. If one is positive, both others are strictly negative; simultaneous zeros of two are impossible.

Theorem 2.2 (The four chambers)**.** *The zero-free model realizes exactly $\mathcal{O}_\star = \{\mathsf{AND}, X, Y, \mathsf{XNOR}\}$. Its fibers and rational interior witnesses are as follows.*

Region	*Truth vector* $(11, 10, 01, 00)$	*Operator*	*Witness* (a, b)
$h_1, h_2, h_3 < 0$	$(1, 0, 0, 0)$	AND	$(1/2, 1/2)$
$h_1 > 0$	$(1, 1, 0, 0)$	X	$(2, 1/2)$
$h_2 > 0$	$(1, 0, 1, 0)$	Y	$(1/2, 2)$
$h_3 > 0$	$(1, 0, 0, 1)$	XNOR	$(3, 3)$

Every fiber is connected. The three discriminant curves are pairwise disjoint in the positive quadrant and are

$$C_1 : \ b = \frac{a}{1+a}, \quad C_2 : \ a = \frac{b}{1+b}, \quad C_3 : \ b = \frac{a}{a-1} \ (a > 1). \tag{3}$$

Each separates AND from exactly one of X, Y, XNOR, respectively, so regular chamber adjacency is a three-leaf star.

Proof. The sign restriction follows from (2), and direct evaluation at the witnesses proves that all four patterns occur. To prove connectedness without a gluing assumption, apply the smooth involution

$$(a, b) \longmapsto (u, v) = (1/a, 1/b). \tag{4}$$

Since $ab > 0$, dividing the scores by ab preserves their signs and gives

$$\frac{1}{ab}(g_{11}, h_1, h_2, h_3) = (u + v + 1,\ v - u - 1,\ u - v - 1,\ 1 - u - v). \tag{5}$$

In the positive (u, v) quadrant the four regions become

X	$v > u + 1$
Y	$u > v + 1$
XNOR	$u + v < 1$
AND	$\lvert u - v\rvert < 1,\ u + v > 1.$

Each is a nonempty intersection of open half-planes, hence is convex and path-connected. Their inverse images under (4) are therefore path-connected. Solving the three score equations gives (3). Equation (2) excludes intersections of two curves in the domain. In reciprocal coordinates each curve is a line segment or ray with nonzero normal, with AND on one side and the corresponding leaf on the other. This also proves the claimed regular adjacency and that exactly one truth-table coordinate changes at each crossing. □

For example, $(a, b) = (2, 1/2)$ gives the scores $(7/2, 1/2, -5/2, -3/2)$ and hence projection X. The output CM is

$$[X] = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}.$$

This is a certified identification, not a conclusion extrapolated from a numerical plot: the witness identifies a region whose uniform signs were proved above.

Remark 2.3 (Positive coefficient rescaling). The same conclusions hold for $F(x, y) = Ax + By + Cxy$ with fixed $A, B, C > 0$. Set $x' = (C/B)x$ and $y' = (C/A)y$; then

$$F(x, y) = \frac{AB}{C}(x' + y' + x'y').$$

The induced positive linear rescaling of magnitude coordinates preserves total polynomial degree, guard-library cardinality, and the operator labels. This extends the example to a positive-coefficient family by an elementary equivalence, not by a new classification argument. The case $C = 0$ is different: the addition model has only the two projection chambers.

3 The polynomial guard-degree tradeoff

3.1 An explicit cost model

A *binary polynomial-sign test* in the original coordinates is

$$p_H(a, b) = \mathbf{1}_{\{H(a,b)>0\}}, \qquad H \in \mathbb{R}[a, b].$$

One test means one such polynomial, not an arbitrary Boolean formula built from several hidden tests. On the promised domain Λ°, a finite library $H_1, \dots, H_k$ is *exact* when its binary signature determines the operator: there is a decoder

$$d : \{0,1\}^k \longrightarrow \mathcal{O}_\star \quad \text{such that} \quad d(p_{H_1}(a,b), \dots, p_{H_k}(a,b)) = \Theta_{(a,b)}$$

for every $(a,b) \in \Lambda^\circ$. Values of d on unrealized signatures do not matter. Tests need not individually be constant on an operator fiber. Constant polynomials can be discarded without improving a minimum. Define m_d to be the minimum k for tests of total degree at most the integer $d \geq 1$, and set $m_d = \infty$ when no finite library is exact.

This cost counts fixed library size, not arithmetic operations, decision-tree depth, coefficient bit complexity, or the cost of proving that an input is zero-free. Strict-sign dispatch on an unpromised domain requires the additional treatment in Section 4.

3.2 The boundary obstruction

Lemma 3.1 (Boundary factors)**.** *The three quadratics h_1, h_2, h_3 in (1) are pairwise nonassociate irreducibles in $\mathbb{R}[a,b]$. If finitely many nonzero polynomial zero sets cover C_i, at least one of their polynomials is divisible by h_i. After constant tests are discarded, every exact finite polynomial-sign library has such a factor for each of $i = 1, 2, 3$.*

Proof. Write $h_i = A_i(a)b + B_i(a)$. The coefficient pairs are

$$(A_1, B_1) = (-a-1, a), \quad (A_2, B_2) = (1-a, -a), \quad (A_3, B_3) = (a-1, -a).$$

In each pair the entries are coprime in $\mathbb{R}[a]$. Thus h_i is primitive and is linear, hence irreducible, over the fraction field $\mathbb{R}(a)$. Gauss's lemma gives irreducibility in $\mathbb{R}[a,b]$. Comparing coefficients shows that no two are scalar multiples.

Parametrize the corresponding positive arc by $b = -B_i(a)/A_i(a)$ on an open interval avoiding the pole. Substituting into any polynomial $H(a,b)$ gives a rational function of a. Unless it vanishes identically, it has only finitely many zeros on that interval. A finite collection of such finite sets cannot cover the interval. Therefore some H vanishes identically after substitution. Division in $\mathbb{R}(a)[b]$, followed by primitivity, shows that h_i divides H in $\mathbb{R}[a,b]$.

Finally, let $z \in C_i$. If every library polynomial were nonzero at z, all test values would be constant on a sufficiently small neighborhood of z. That neighborhood contains zero-free points from the two different operator chambers adjoining C_i, contradicting exactness. Hence the library zero sets cover C_i, and the preceding argument supplies the required factor. This reasoning uses points on the boundary only to compare nearby promised inputs; it does not assume that dispatch is defined on Δ. □

Theorem 3.2 (Exact degree–cardinality tradeoff)**.** *For the four-operator model in the original magnitude coordinates,*

$$m_d = \begin{cases} \infty, & d = 1, \\ 3, & d = 2 \text{ or } 3, \\ 2, & d \geq 4. \end{cases} \tag{6}$$

Furthermore, in any exact library consisting of two polynomial-sign tests, each *defining polynomial has degree at least four.*

Proof. A polynomial of degree at most one cannot contain any of the quadratic factors required by Lemma 3.1. Consequently no finite affine-sign library suffices. The same argument rules out any finite Boolean combination of affine inequalities as a description of the complete operator atlas: all such combinations over a finite library still depend only on its signature.

A polynomial of degree at most three cannot contain two distinct h_i, since their product has degree four and they are relatively prime. Covering all three boundaries therefore requires at least three tests. The three tests $\mathbf{1}_{\{h_i>0\}}$ achieve this bound by Theorem 2.2.

For degree four, take the two tests

$$r_1 = \mathbf{1}_{\{h_1 h_3<0\}}, \qquad r_2 = \mathbf{1}_{\{h_2 h_3<0\}}. \tag{7}$$

They have the required $H > 0$ form with $H = -h_1 h_3$ and $H = -h_2 h_3$. Their signatures are

Operator	r_1	r_2
AND	0	0
X	1	0
Y	0	1
XNOR	1	1

Thus two tests suffice. One binary test cannot distinguish four nonempty operator fibers, so two is optimal.

For the stronger assertion, an exact two-test library has only four possible signatures. Since four operator labels occur and signatures cannot be shared by different labels, each operator must have exactly one signature and all four signatures are used. Relative to the AND signature, each of the two bits changes on exactly two of the other three signatures. The polynomial defining that bit must consequently vanish along the two corresponding AND/leaf boundaries. Applying the substitution argument in Lemma 3.1 to that polynomial gives two distinct quadratic factors, so its degree is at least four. The argument applies to both tests. □

The lower bound covers arbitrary real coefficients and arbitrary decoders; it does not merely check a small catalogue of candidate polynomials. The accompanying code checks the displayed factorizations and all four output signatures, but the universal minimum is the consequence of the proof, not of a numerical search.

3.3 What the separation does not imply

The affine obstruction is not invariant under arbitrary changes of coordinates. Equations (4)–(5) give an explicit smooth change under which the three original score signs are affine tests. Alternatively, introducing the feature $z = ab$ makes them linear in (a, b, z) on the feature surface $z = ab$. Either transformation changes the permitted input features and therefore the observation contract.

Nor does two-test quartic dispatch establish a runtime improvement. Computing $h_1 h_3$ and $h_2 h_3$ from the original scores uses those scores and two additional multiplications. The theorem compares sign-information channels under a degree constraint; feature construction, arithmetic sharing, numerical certification, and bit complexity must be charged separately in an implementation study.

4 Zero-aware certification and partial abstraction

4.1 Nonempty guards and explicit coverage

A strict sign certificate on a guard $G \subseteq \Lambda$ proves either $g_\alpha > 0$ throughout G or $g_\alpha < 0$ throughout G. It may be an analytic proof, an exact real-algebraic certificate, or another sound verification object. A sample point or a plot is not a certificate of a uniform sign.

Proposition 4.1 (Nonempty-guard soundness)**.** *Let G be nonempty. If one strict sign of every g_α is proved throughout G, then $G \subseteq \Lambda^\circ$ and G carries a unique exact Boolean operator and correspondence tensor. A collection of such guards is exact on a declared domain $D_0 \subseteq \Lambda^\circ$ when it covers D_0. Guards with the same operator may be merged by explicit union; overlapping nonempty guards with certified labels cannot disagree on their intersection.*

Proof. Strict signs exclude zeros and determine the truth value at every input. Choosing any point of G proves that this truth table is realized, and the uniform certificates make it the same at every point. Coverage and the assertions about unions and intersections follow pointwise. □

Empty guards must be discarded or marked infeasible. All strict inequalities are vacuously true there, so the operator is not uniquely determined by empty-guard semantics. This is why nonemptiness is part of the statement rather than an implicit implementation preference.

Polynomial sign decomposition and quantifier elimination provide exact routes in principle [4, 16]. Interval methods and solver-based methods may leave cells uncertified. Such an outcome reports a limit of the analysis; it is not evidence that both signs actually occur. The present note supplies no general solver for arbitrary continuous scores.

4.2 Certified dictionaries and reduced observation

A complete, already proved operator dictionary can justify inference of unobserved signs.

Proposition 4.2 (Reduced certificate with an explicit promise)**.** *Let $D_0 \subseteq \Lambda^\circ$ be a certified zero-free domain and let the finite set $\mathcal{O}$ contain every Θ_λ for $\lambda \in D_0$. Let $G \subseteq D_0$ be nonempty. Suppose strict signs on a selected set $S \subseteq \{0,1\}^n$ are proved constant throughout G, and exactly one member $\Theta \in \mathcal{O}$ has the corresponding restriction to S. Then every point of G has operator Θ, including the truth values at unobserved inputs.*

Proof. At any point of G an induced total operator exists by the zero-free hypothesis, lies in $\mathcal{O}$ by dictionary completeness, and agrees with the certified restriction. Uniqueness of the dictionary match identifies it with Θ. □

Neither dictionary completeness nor the zero-free promise may be inferred from a sample. Without the latter, even an empty separating test set for a singleton dictionary can be unsafe: the scores 1 and t^2 induce the constant-true operator for every $t \neq 0$, but not at $t = 0$.

Boundary rejection in the nonlinear example. At $(a, b) = (1, 1/2)$,

$$(g_{11}, h_1, h_2, h_3) = (2, 0, -1, -1).$$

All three binary predicates $h_i > 0$ are false, exactly as in the AND chamber; so they cannot by themselves certify membership in Λ°. A correct unpromised dispatcher uses exact trichotomy,

rejecting a zero rather than silently treating it as negative. The quartic method likewise checks whether either product in (7) is zero. Since $g_{11} > 0$, both products are nonzero exactly when all h_i are nonzero. Thus two *ternary* product-sign observations suffice for classification together with boundary rejection, but this is a different outcome alphabet from the binary contract in Theorem 3.2.

4.3 Collecting semantics for predicate cells

Synthetic representatives supply one score per Boolean input. Predicate abstraction of a concrete state space need not do so. Let D be a state space, $P : D \to \{0,1\}^n$ a total predicate map, and $F : D \to \mathbb{R}$ a total score. Define

$$S_\alpha = \{\operatorname{sgn}_0 F(x) : x \in D,\ P(x) = \alpha\} \subseteq \{-1, 0, +1\},$$

where sgn_0 retains zero. The empty set means an unreachable cell. The singleton sets $\{-1\}$ and $\{+1\}$ mean exact false and true outputs. A set containing 0 exhibits a tie, and a set containing both strict signs exhibits genuine two-sided variation. The last two properties can hold simultaneously.

Proposition 4.3 (Exactness on predicate cells)**.** *On a reachable cell, $q \circ F$ is everywhere defined and constant exactly when S_α is one of $\{-1\}, \{+1\}$. Under the declared full-cube convention, a canonical total Boolean operator is induced exactly when every cell is reachable and satisfies this singleton criterion. Otherwise one retains a reachability mask and exact/partial cell information, unless a separate completion policy is declared.*

Proof. The two singleton cases make the threshold defined and constant. Conversely, a defined constant strict threshold permits neither a zero nor both strict signs. Values on unreachable inputs are not fixed by the concrete model; adding them is a completion choice. □

This is a collecting-semantics contract, not a new abstract domain [9]. In particular, checking only existence of a positive score and existence of a negative score cannot distinguish an empty cell from a nonempty all-zero cell: both tests are false in both cases. Reachability, zeros, genuine variation, and a solver's `UNKNOWN` response must remain separate.

5 Test selection, coordinate probes, and faithful representations

5.1 Finite predicate selection is a minimum-test-set problem

Suppose a certified finite partition $\{A_1, \dots, A_m\}$ of a zero-free domain has nonempty atoms labelled by operators Θ_i. Let $p_1, \dots, p_r$ be admissible binary predicates, each constant on each atom, and let $w_j > 0$ be selection costs. Put

$$D_{ik} = \{j : p_j(A_i) \neq p_j(A_k)\} \quad (\Theta_i \neq \Theta_k).$$

A minimum-cost exact *fixed predicate basis* is the binary optimization problem

$$\begin{aligned} &\min \sum_{j=1}^{r} w_j z_j, \\ &\sum_{j \in D_{ik}} z_j \geq 1 \quad (\Theta_i \neq \Theta_k), \qquad z_j \in \{0,1\}. \end{aligned} \tag{8}$$

Indeed, selected predicates yield exact labels if and only if they separate every differently labelled pair, which is exactly the set of displayed constraints. If any required D_{ik} is empty, no exact basis exists in that library. If predicates vary within an atom, the partition must first be refined and certified; testing only its representatives does not establish this premise.

This is the classical pair-separation/hitting-set formulation [17, 6]. It does not optimize expected path length or adaptive worst-case depth. Its weights account for selected predicates only; costs of computing shared features are not automatically additive in these weights. Predicate abstraction with minimum predicates [8] is an important related precedent, but exact label recovery here is a different specification from preserving a chosen verification property.

5.2 Fixed coordinates versus adaptive and target-specific information

For a nonempty finite family $\mathcal{O}$ of Boolean functions, define

$$\sigma_{\mathrm{CM}}(\mathcal{O}) = \min\{|S| : S \subseteq \{0,1\}^n,\ \Theta \mapsto \Theta|_S \text{ is injective on } \mathcal{O}\}.$$

It is zero for a singleton. This is a fixed truth-table-coordinate test-set number, expressed in CM notation, not a new generic combinatorial invariant. Its exact formulation is (8) with one test per coordinate and unit cost; hence

$$\lceil \log_2 |\mathcal{O}| \rceil \le \sigma_{\mathrm{CM}}(\mathcal{O}) \le 2^n.$$

The lower bound counts binary signatures and the upper bound uses the complete table.

It must not be confused with a target-specific specification or teaching set, whose chosen coordinates can depend on the particular function being specified, or with adaptive querying of an unknown function. Threshold-function specification numbers and their chamber-facet interpretation are developed in [2]; the global, fixed separating set above has different quantifiers.

Proposition 5.1 (The coordinate-oracle costs of the four-operator family)**.** *For $\mathcal{O}_\star$, the unique minimum fixed coordinate set is $\{10, 01, 00\}$, so $\sigma_{\mathrm{CM}}(\mathcal{O}_\star) = 3$. Deterministic adaptive identification from truth-table-coordinate queries also has worst-case cost three. AND requires three coordinate values for a target-specific certificate, while each leaf operator has a one-coordinate certificate.*

Proof. Coordinate 11 is always 1, and the other coordinates are

	10	01	00
AND	0	0	0
X	1	0	0
Y	0	1	0
XNOR	0	0	1.

Omitting a nonconstant coordinate leaves AND indistinguishable from the leaf differing only there. This proves the fixed minimum and its uniqueness. On the AND input, an adaptive algorithm receives zero at every nonconstant coordinate. Until all three have been queried, the leaf at an unqueried coordinate is still consistent with the transcript; the worst-case lower bound follows. Querying all three proves the upper bound. A positive value at the unique leaf coordinate identifies that leaf, whereas AND needs all three zeros. □

The adaptive lower bound is for an unknown operator accessible only through the listed coordinate oracle and a known dictionary. An algorithm given numerical (a, b) and allowed arbitrary derived

predicates has a different contract, as the quartic construction illustrates. Within the coordinate oracle, an early-stop order has depths $1, 2, 3, 3$ for the four operators and expected cost $9/4$ under the uniform distribution. Thus "three probes are necessary" means a fixed-basis or worst-case claim, not a claim about every parameter point.

5.3 The fixed-frame LM endpoint is self-contained

Let $\mathcal{X} = (X_1, \ldots, X_n)$ be independent formal Boolean generators, and identify formulas up to Boolean equivalence. For a formula A, set $A^{\langle 1 \rangle} = A$ and $A^{\langle 0 \rangle} = \neg A$. Extend a Boolean function Θ by its disjoint minterms:

$$\widetilde{\Theta}(A_1, \ldots, A_n) = \bigvee_{\beta \in \{0,1\}^n} \left(\Theta(\beta) \wedge \bigwedge_{i=1}^{n} A_i^{\langle \beta_i \rangle} \right).$$

The minterms are disjoint, so an iterated XOR, denoted $\Updownarrow$, would give the same function. Define the formula-valued lift

$$L_{\mathcal{X}}(\Theta)[\alpha] = \widetilde{\Theta}(X_1^{\langle \alpha_1 \rangle}, \ldots, X_n^{\langle \alpha_n \rangle}). \tag{9}$$

The formal reference frame is not a collection of constrained physical predicates: all Boolean valuations, including the all-true valuation, are available.

Proposition 5.2 (Faithfulness and representation-invariant dispatch)**.** *All-true valuation of* (9) *returns* $[\Theta]$*. Consequently* $L_{\mathcal{X}}$ *is injective. For every Boolean connective* $H : \{0,1\}^k \to \{0,1\}$*, the lift preserves pointwise superposition:*

$$L_{\mathcal{X}}(H(\Theta_1, \ldots, \Theta_k)) = H(L_{\mathcal{X}}(\Theta_1), \ldots, L_{\mathcal{X}}(\Theta_k)).$$

Therefore

$$\Theta_\lambda = \Theta_\mu \iff [\Theta_\lambda] = [\Theta_\mu] \iff L_{\mathcal{X}}(\Theta_\lambda) = L_{\mathcal{X}}(\Theta_\mu).$$

Any fixed parameter-side observation contract and cost has the same exact-dispatch optimum for these three outputs, provided decoding or output-construction costs are not added differently.

Proof. At $X_i = 1$, the literal $X_i^{\langle \alpha_i \rangle}$ has value α_i; the minterm extension therefore evaluates to $\Theta(\alpha)$. Equal lifted tensors give equal truth tables under this valuation, proving injectivity. At every entry and every valuation of the formal frame, both sides of the superposition identity evaluate the same connective H on the same function values. Finally, injective representations induce the same equality relation on parameter space. The same observations consequently identify the same labels; only representational output costs could alter a total implementation cost. □

For the running point $(2, 1/2)$, the complete output chain is

$$(2, 1/2) \longmapsto X \longmapsto \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix} \longmapsto \begin{bmatrix} X_1 & X_1 \\ \neg X_1 & \neg X_1 \end{bmatrix}.$$

For comparison, the AND and XNOR lifts are

$$\begin{bmatrix} X_1 \wedge X_2 & X_1 \wedge \neg X_2 \\ \neg X_1 \wedge X_2 & \neg X_1 \wedge \neg X_2 \end{bmatrix}, \qquad \begin{bmatrix} X_1 \Leftrightarrow X_2 & X_1 \Updownarrow X_2 \\ X_1 \Updownarrow X_2 & X_1 \Leftrightarrow X_2 \end{bmatrix}.$$

Continuous scores, numeric truth tensors, and formula-valued lifts are different types. A partial truth tensor from an unreachable or nonexact predicate cell does not determine a full lift without a completion policy. The elementary proof above removes any dependence on an unpublished companion theorem; the companion is a terminology reference only.

6 Reproducibility, limitations, and conclusion

The companion archive contains the source of this note, a portable exact rational dispatcher, a deterministic verification script, machine-readable results, reproduction commands, and an audit crosswalk. Both coordinate and quartic dispatch modes return a separate boundary status; floating-point inputs are rejected by the reference implementation to avoid silently replacing exact rational semantics with approximate comparisons. The tests compare direct signed-input evaluation, the expanded scores, the reciprocal normal form, and both dispatch modes. They also check boundary points, the original recorded rational grid, dictionary separation, actual minterm evaluation, and the classical Walsh witnesses.

These checks are bounded computations, not formal proof certificates for the topology or universal polynomial lower bound. The latter rest on the proofs in Theorems 2.2 and 3.2. The package documents exactly what was run; it does not claim solver benchmarks, a speedup over direct score evaluation, or a completed general guard-synthesis implementation.

The example separates three issues that can otherwise be conflated. Geometry determines where each operator is realized. An observation contract determines which tests may be used and how they are counted. Certification determines when those observations justify a total Boolean output. In the original coordinates, the exact degree–cardinality tradeoff (6) and the three-coordinate lower bound make the contract dependence explicit. Reciprocal linearization prevents interpreting the result as a coordinate-free nonlinear obstruction. Faithful CM/LM representations preserve dispatch distinctions but do not strengthen the information available from the continuous model.

The general quotient, collecting semantics, finite test selection, and Walsh arrangement remain classical or elementary components. A stronger algorithms or applications claim would require a motivated model family, an implemented certificate-producing backend, and an evaluation that charges feature construction and certification as well as sign queries. Those are not conclusions of this note.

A Classical trajectory and Walsh calibrations

A.1 Crossings, event counts, and endpoint-relative intervals

Proposition A.1 (Single-entry crossing)**.** *Let $\gamma : (-\eta, \eta) \to \Lambda$ be continuous. Suppose one score has a fixed strict sign on $(-\delta, 0)$ and the opposite fixed strict sign on $(0, \delta)$, while all other scores are nonzero on $(-\delta, \delta)$. Then the two side operators differ at exactly that entry. If the composed scores are differentiable, a simple zero of one score and nonzero values of all the others at 0 suffice after shrinking δ.*

Proof. Every other composed score keeps its sign by continuity. The distinguished entry alone changes. A nonzero derivative at a zero gives opposite strict signs on sufficiently short punctured sides. □

A mere assertion that a continuous function "changes sign at zero" would not exclude oscillation. The fixed-side hypotheses are therefore part of the statement. One-input adjacency is already a feature of structured parameter graphs and threshold arrangements [11, 2].

Proposition A.2 (Trajectory and polynomial event bound)**.** *For a continuous path $\gamma : [0, T] \to \Lambda$, the set $S = \gamma^{-1}(\Delta)$ is closed and the induced operator is constant on each connected component of $[0, T] \setminus S$. If all points of S are isolated relative to $[0, T]$, then S is finite. The complementary*

components are relatively open *intervals, possibly containing* 0 *or* T*. Separately, if the scores on an interval are nonzero polynomials of degree at most* d*, there are at most* $2^n d$ *discriminant values and at most* $2^n d + 1$ *complementary interval components.*

Proof. Relative closedness of Δ and continuity give closedness of S. Connected complementary components map into one fiber by Proposition 2.1. An infinite closed subset of a compact interval has an accumulation point in the set, contrary to isolation. For the polynomial assertion, each nonzero polynomial has at most d roots and the union bound gives the result. Even-multiplicity roots need not change the operator, so this counts possible zero events, not necessarily actual switches. □

These are elementary score-level observations, not a dynamics theorem or a substitute for logical bifurcation theory [1].

A.2 A centered Walsh arrangement

Let $N = 2^n$ and $x \in \{-1, +1\}^n$. For nonempty $S \subseteq [n]$, put $\chi_S(x) = \prod_{i \in S} x_i$ and let $\phi(x)$ be the vector of these $N - 1$ characters. Standard Walsh orthogonality [18] gives

$$\phi(x) \cdot \phi(y) = \begin{cases} N - 1, & x = y, \\ -1, & x \neq y, \end{cases} \qquad \sum_x \phi(x) = 0.$$

The coefficient family $F_c(x) = c \cdot \phi(x)$ identifies $\mathbb{R}^{N-1}$ with the zero-sum subspace of score vectors in $\mathbb{R}^N$: the displayed Gram matrix is $NI - J$, which has rank $N - 1$.

Proposition A.3 (Classical Walsh calibration)**.** *For every* $n \geq 1$*, precisely the* $2^N - 2$ *nonconstant strict Boolean sign patterns occur, each in one open convex cone. For* $n \geq 2$*, regular facet adjacency is Hamming-distance-one adjacency on those truth tables, the* N*-cube with the two constant vertices deleted. For* $n = 1$ *the two indexed zero hyperplanes coincide and ordinary geometric adjacency is not single-entry adjacency.*

Proof. For a positive set P of size $0 < k < N$, choose $c = \sum_{y \in P} \phi(y)$. Then

$$c \cdot \phi(x) = \begin{cases} N - k, & x \in P, \\ -k, & x \notin P. \end{cases}$$

This realizes every nonconstant pattern. The zero-sum identity excludes the two constant strict patterns. Every fiber is an intersection of open half-spaces, hence convex.

For $n \geq 2$, consider two realized sign vectors differing only at coordinate j. The common signs on the other $N - 1$ coordinates must include both signs; otherwise one of the two full patterns would be constant. Choose nonzero values with those signs whose sum is zero, and set coordinate j to zero. This is a point on the single indexed hyperplane, with no other coordinate zero. A small perturbation in the zero-sum subspace to either side realizes the two patterns, proving regular adjacency. Conversely, at a regular single-hyperplane crossing only that coordinate changes. For $n = 1$, the score vector has the form $(t, -t)$; the two rays meet at a simultaneous two-coordinate zero, giving the exception. □

The count and tope graph are classical homogeneous-dichotomy and arrangement geometry [10, 7]. This appendix supplies a directly checkable calibration of the interface, not an independent novelty claim for Walsh expansions or chamber counts.

References

[1] W. Abou-Jaoudé and P. T. Monteiro, On logical bifurcation diagrams, *Journal of Theoretical Biology* 466 (2019), 39–63. doi:10.1016/j.jtbi.2019.01.008.

[2] M. Anthony, Chamber geometry and specification numbers of Boolean threshold functions, arXiv:2606.29477 (2026). `https://arxiv.org/abs/2606.29477`.

[3] G. Averkov and L. Bröcker, Minimal polynomial descriptions of polyhedra and special semialgebraic sets, arXiv:1002.0921 (2010). `https://arxiv.org/abs/1002.0921`.

[4] S. Basu, R. Pollack, and M.-F. Roy, *Algorithms in Real Algebraic Geometry*, 2nd ed., Springer, 2006.

[5] S. Basu, R. Pollack, and M.-F. Roy, An asymptotically tight bound on the number of semialgebraically connected components of realizable sign conditions, *Combinatorica* 29 (2009), 523–546. doi:10.1007/s00493-009-2357-x.

[6] P. Berman, B. DasGupta, and M.-Y. Kao, Tight approximability results for test set problems in bioinformatics, *Journal of Computer and System Sciences* 71(2) (2005), 145–162. doi:10.1016/j.jcss.2005.02.001.

[7] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. M. Ziegler, *Oriented Matroids*, 2nd ed., Cambridge University Press, 1999.

[8] S. Chaki, E. Clarke, A. Groce, and O. Strichman, Predicate abstraction with minimum predicates, in *CHARME 2003*, LNCS 2860, Springer, 2003, pp. 19–34. doi:10.1007/978-3-540-39724-3_5.

[9] P. Cousot and R. Cousot, Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints, in *POPL 1977*, pp. 238–252. doi:10.1145/512950.512973.

[10] T. M. Cover, Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition, *IEEE Transactions on Electronic Computers* EC-14(3) (1965), 326–334. doi:10.1109/PGEC.1965.264137.

[11] P. Crawford-Kahrl, B. Cummins, and T. Gedeon, Joint realizability of monotone Boolean functions, *Theoretical Computer Science* 922 (2022), 447–474. doi:10.1016/j.tcs.2022.04.045.

[12] B. Cummins, T. Gedeon, S. Harker, K. Mischaikow, and K. Mok, Combinatorial representation of parameter space for switching networks, *SIAM Journal on Applied Dynamical Systems* 15(4) (2016), 2176–2212. doi:10.1137/15M1052743.

[13] B. Cummins, T. Gedeon, S. Harker, and K. Mischaikow, DSGRN: Examining the dynamics of families of logical models, *Frontiers in Physiology* 9 (2018), 549. doi:10.3389/fphys.2018.00549.

[14] B. Cummins, M. Gameiro, T. Gedeon, K. Mischaikow, and B. Rivas, Boolean models coarsely sample continuous dynamics of regulatory networks, arXiv:2606.14925 (2026). `https://arxiv.org/abs/2606.14925`.

[15] T. Gedeon, Multi-parameter exploration of dynamics of regulatory networks, *BioSystems* 190 (2020), 104113. doi:10.1016/j.biosystems.2020.104113.

[16] G. Jeronimo, D. Perrucci, and J. Sabia, On sign conditions over real multivariate polynomials, *Discrete & Computational Geometry* 44 (2010), 195–222. doi:10.1007/s00454-009-9200-4.

[17] B. M. E. Moret and H. D. Shapiro, On minimizing a set of tests, *SIAM Journal on Scientific and Statistical Computing* 6(4) (1985), 983–1003. doi:10.1137/0906067.

[18] R. O'Donnell, *Analysis of Boolean Functions*, Cambridge University Press, 2014.

[19] B. Theory, Correspondence and Logical Matrices: A Boolean Operator Calculus, unpublished companion manuscript, 2026. Cited for terminology; no proof in this note depends on it.