Entropy Dispersion and Total Influence: A Structural Route Toward Conditional P ≠ NP

Keith Taylor VERSF Theoretical Physics Program

General Reader Summary

What this paper is about. Some problems are easy to check but seemingly hard to solve. You can verify a completed Sudoku in seconds, but finding the solution from scratch can take far longer. The deepest open question in mathematics and computer science — the P versus NP problem — asks whether this gap between checking and finding is fundamental or merely a reflection of our current ignorance.

What the parent paper showed. A companion manuscript argues that the gap is fundamental, rooted in the same physical principles that govern heat flow and irreversibility. It proves that two major families of algorithms — logical deduction (used by industrial problem-solvers) and statistical detection (used by machine learning) — provably cannot solve hard instances efficiently. The remaining question is whether some third approach could succeed.

What this paper does. We provide the mathematical tools needed to analyze that remaining question rigorously. Our central idea is *algorithmic total influence*: a way to measure how sensitive an algorithm's success is to the specific details of the problem it is trying to solve. We prove three things:

1. **If an algorithm succeeds, it must be paying attention.** Any algorithm that beats random guessing must be sensitive to at least some specific features of the problem instance. We quantify exactly how much sensitivity is required relative to the advantage gained (Section 2).
2. **Attention accumulates one step at a time.** An algorithm builds its advantage through a sequence of small information-gathering steps, like a detective collecting clues. We prove that the total advantage is bounded by the sum of what each step contributes, and that each step's contribution is limited by how much of the problem it can see (Section 3).
3. **For most known types of algorithms, sensitivity leaves a detectable trace.** We show that for algorithms with limited memory, algorithms that inspect the problem through local queries, and algorithms built from simple logical circuits, the required sensitivity takes a specific mathematical form — a "footprint" — that falls within the scope of known impossibility results (Section 4).

What remains open. The one gap we cannot close is whether *every* efficient algorithm must leave such a footprint, or whether some exotic, globally coordinated computation could evade detection. We formulate this as a precise conjecture and show that resolving it would complete the argument.

The bottom line. This paper converts physical intuitions about information flow into rigorous mathematical theorems. It does not prove P ≠ NP, but it reduces the problem to a single, sharply defined structural question about the nature of efficient computation.

Table of Contents

References

Abstract

This companion paper formalizes the entropy-dispersion framework underlying the structural conditional separation program for P ≠ NP. We define *algorithmic total influence* with respect to instance encodings and prove a bounded total-influence theorem under local access (Theorem 2.1). We establish that advantage forces an influence spike on at least one accessed coordinate (Theorem 2.2) under explicit baseline-stability conditions that hold for random k-SAT distributions. We prove a martingale-based entropy-dispersion bound on cumulative success bias (Theorem 3.2) with an explicit corollary recovering $\Omega(\varepsilon)$ variation under a neutral prior (Corollary 3.3). We prove Information-to-Structure theorems for bounded-space algorithms (Theorem 4.2, yielding a transcript footprint), local-access algorithms (Theorem 4.4, yielding an influence spike), and AC^0 circuits (Theorem 4.6, yielding a low-degree Fourier footprint). We address the basis-translation problem between instance-encoding coordinates and assignment-variable polynomials (Section 5), producing instance-relative low-degree witnesses and identifying the precise conditions under which these connect to existing lower bounds. Together, these results sharpen the conditional pathway from bounded entropy flow to algorithmic rigidity, proving new theorems for broad computational models and isolating a single, well-defined open problem for the general case.

Keywords: P vs NP, entropy dispersion, total influence, algorithmic hardness, proof complexity, statistical query lower bounds

---

1. Introduction and Motivation

1.1 Context

The parent manuscript (*Operational Separation of Discovery and Verification*) establishes exponential lower bounds for NP search in two complementary algorithmic paradigms — resolution/CDCL proof search and statistical query/low-degree polynomial estimation — and identifies the Global-to-Local Reduction (GLR) conjecture as the single remaining obstruction to a complete operational separation.

This companion paper addresses three specific gaps identified in critical review of the parent manuscript:

1. **The emergent-time-to-locality gap.** The parent paper argues that emergent time forces incremental information acquisition, but the step from "incremental" to "bounded-local in the instance" requires additional mathematical structure. We formalize this via algorithmic total influence and prove that local-access algorithms have bounded total

influence (Section 2), then establish that bounded total influence constrains the rate of entropy dispersion (Section 3).
2. **The structure-extraction gap.** The parent paper's Information-to-Structure claim (connecting detectable advantage to resolution or low-degree footprints) was stated at a level of generality that outran the available proof. We prove this rigorously for bounded-space, local-access, and AC^0 models (Section 4), with each theorem stating its precise conclusion. The general case is stated as a conjecture with explicit falsification conditions.
3. **The basis-mismatch problem.** The parent paper's k-local predicates are defined over instance-encoding coordinates, while the SQ/low-degree lower bounds concern polynomials over assignment variables. We prove a basis-translation theorem (Section 5) producing instance-relative low-degree witnesses, and identify the precise distributional conditions under which this connects to existing lower bounds.

1.2 Notation and Conventions

Throughout, we use the following notation:

- Φ denotes an NP instance with witness length n
- $E(\Phi) \in \{0,1\}^M$ denotes the encoding of Φ, where $M = M(n) = \mathrm{poly}(n)$
- $V(\Phi, x) \in \{0,1\}$ is the polynomial-time verifier
- $S\_\Phi := \{x \in \{0,1\}^n : V(\Phi, x) = 1\}$ is the solution set
- $p_0(\Phi) := |S\_\Phi| / 2^n$ is the baseline success probability
- $\mathcal{D}_n$ denotes a distribution over instances of witness length n
- A denotes a randomized algorithm with internal randomness r
- $S := \mathbb{1}[V(\Phi, A(\Phi; r)) = 1]$ is the success indicator
- $p\_A(\Phi) := \Pr_r[S = 1 \mid \Phi]$ is the algorithm's success probability on instance Φ
- negl(n) denotes any function that is $o(n^{-c})$ for all constants $c > 0$

Distributional Convention. Where results require specific distributional properties, these are stated explicitly as hypotheses. Theorem 2.1, Theorem 3.2, and Theorem 3.5 hold for arbitrary distributions $\mathcal{D}_n$. Theorem 2.2 requires baseline stability (H1)–(H2). Corollary 3.3 requires three explicit conditions (C1)–(C3). Theorem 4.6 requires product structure for the Fourier analysis.

---

2. Algorithmic Total Influence

2.1 Definitions

We formalize the sensitivity of solver success to local perturbations of the instance encoding.

Definition 2.1 (Instance-Bit Influence). For each encoding coordinate $j \in [M]$, define the *instance-bit influence* of coordinate j on algorithm A under distribution $\mathcal{D}$:

$\mathrm{Inf}_j(A; \mathcal{D}) := \mathbb{E}\_\{\Phi \sim \mathcal{D}\}[\ |p\_A(\Phi) - p\_A(\Phi^\{\oplus j\})|\ ]$

where $\Phi^\{\oplus j\}$ denotes the instance obtained by flipping bit j of $E(\Phi)$.

Remark 2.1a (Flip-Closure). When $\mathcal{D}$ is not closed under bit flips (e.g., random k-SAT conditioned on satisfiability), we interpret $\Phi^\{\oplus j\}$ via the natural resampling coupling on the underlying product generative model, or we restrict attention to the flip-closed support. All theorems that evaluate $p\_A(\Phi^\{\oplus j\})$ explicitly state the required distributional setting. In particular, Theorem 2.1 is valid for any $\mathcal{D}$ because its proof uses only a coupling argument that does not require $\Phi^\{\oplus j\}$ to lie in the support, while Theorem 2.2 assumes baseline stability (H2) which ensures that p_0 is well-behaved under single-coordinate flips.

Definition 2.2 (Total Influence). The *total influence* of algorithm A under distribution $\mathcal{D}$ is:

$\mathrm{TInf}(A; \mathcal{D}) := \Sigma_{j=1}^{M} \mathrm{Inf}_j(A; \mathcal{D})$

Definition 2.3 (Maximum Influence). The *maximum influence* is:

$\mathrm{MaxInf}(A; \mathcal{D}) := \max_j \in [M]\ \mathrm{Inf}_j(A; \mathcal{D})$

Remark 2.4. Total influence measures the aggregate sensitivity of solver success to all local instance perturbations. An algorithm with low total influence is one whose success probability is robust to local changes in the instance. An algorithm with high total influence necessarily depends on specific encoding bits, creating a detectable structural footprint.

2.2 Local-Access Total-Influence Bound

Theorem 2.1 (Local-Access Influence Bound). Let A be a randomized polynomial-time algorithm that interacts with instance Φ via at most $t = t(n)$ adaptive queries, where each query inspects at most $q = q(n)$ bits of the encoding $E(\Phi)$. Then for any distribution $\mathcal{D}$:

$\mathrm{TInf}(A; \mathcal{D}) \leq 2 \cdot t(n) \cdot q(n)$

Proof.

Let $Q_1, Q_2, \ldots, Q_t$ be the (random, adaptive) queries issued by A during execution on instance Φ with randomness r. Each query Q_i inspects a set of coordinates $J_i \subseteq [M]$ with $|J_i| \leq q$.

Define the *queried set* $J := J_1 \cup J_2 \cup \cdots \cup J_t$. Since $|J_i| \leq q$ for each i, we have $|J| \leq t \cdot q$.

Fix any coordinate $j \in [M]$. Consider the coupling between executions of A on Φ and on $\Phi^\{\oplus j\}$, using the same randomness r. The transcripts diverge only if some query Q_i reads coordinate j. Let $B_j(\Phi, r) := \mathbb{1}[j \in J(\Phi, r)]$ be the event that j is queried. Then:

$|p\_A(\Phi) - p\_A(\Phi^\{\oplus j\})| \leq \Pr_r[B_j(\Phi, r) = 1 \mid \Phi] + \Pr_r[B_j(\Phi^\{\oplus j\}, r) = 1 \mid \Phi^\{\oplus j\}]$

This follows because on executions where j is never read, the outputs are identically distributed; divergence can only occur when j is read by at least one query.

Averaging over Φ and summing over all j:

$\mathrm{TInf}(A; \mathcal{D}) = \Sigma_j \mathbb{E}\_\Phi[|p\_A(\Phi) - p\_A(\Phi\text{^}\{\oplus j\})|] \leq 2 \cdot \mathbb{E}\_\Phi[\Sigma_j \Pr_r[j \in J(\Phi, r) \mid \Phi]] = 2 \cdot \mathbb{E}\_\{\Phi,r\}[|J(\Phi, r)|] \leq 2 \cdot t \cdot q$ ■

Remark 2.5. This bound holds for *any* distribution $\mathcal{D}$ because it depends only on the algorithm's query structure, not on the distribution. No product, flip-closure, or spectral gap assumption is needed.

Remark 2.6. The bound is tight. An algorithm that reads $t \cdot q$ distinct bits and outputs their XOR as a candidate witness bit achieves $\Omega(1)$ influence on each queried coordinate.

2.3 Converse: Advantage Forces an Influence Spike

Theorem 2.1 bounds total influence from above under local access. The following result establishes the converse: if a local-access algorithm achieves advantage, some accessed coordinate must carry non-negligible influence. This is the strongest universally correct converse in our setting — it requires no product distribution, no hypercube Poincaré inequality, and no flip-closure of $\mathcal{D}$.

Theorem 2.2 (Influence Spike Under Baseline Stability). Let A be a randomized algorithm that, on input Φ, makes at most $t = t(n)$ adaptive queries, each inspecting at most $q = q(n)$ encoding bits. Let

$\varepsilon := \mathbb{E}\_\{\Phi \sim \mathcal{D}\}[p\_A(\Phi) - p_0(\Phi)] \geq 0$

Assume:

(H1) **Baseline domination:** $p\_A(\Phi) \geq p_0(\Phi)$ for all Φ in the support of $\mathcal{D}$. (Can always be enforced by mixing in uniform guessing with small probability.)

(H2) **Baseline stability:** The baseline success probability $p_0(\Phi)$ is stable under single-coordinate changes: for any Φ and any $j \in [M]$, $|p_0(\Phi) - p_0(\Phi\text{^}\{\oplus j\})| \leq \mathrm{negl}(n)$.

Then there exists a query step $i \in \{1, \ldots, t\}$ and an encoding coordinate j inspected by the i-th query such that:

$\mathrm{Inf}_j(A; \mathcal{D}) \geq \Omega(\varepsilon / (tq))$

Remark (When the Hypotheses Hold). For random k-SAT at the satisfiability threshold, (H2) holds because flipping a single clause-encoding bit changes the number of solutions by at most a $(1 \pm O(1/n))$ factor, so p_0 shifts by at most negl(n). Condition (H1) is enforceable without loss of generality. These conditions are satisfied in all settings relevant to the parent manuscript, but they are genuine hypotheses — the theorem does not hold for arbitrary distributions where p_0 is highly sensitive to single-bit changes.

Proof.

Let $J(\Phi, r) \subseteq [M]$ be the set of coordinates inspected by A on instance Φ with randomness r, so $|J(\Phi, r)| \leq tq$.

The algorithm's success probability $p\_A(\Phi)$ is determined by $E(\Phi)$ and the distribution over r. The only channel through which instance information enters the computation is the queried coordinates $J(\Phi, r)$. Therefore, for any coordinate $j \notin J(\Phi, r)$, flipping bit j does not affect the output distribution, and hence:

If $j \notin J(\Phi, r)$ for all r, then $p\_A(\Phi) = p\_A(\Phi^\{\oplus j\})$

Now suppose for contradiction that every coordinate j satisfies:

$\mathrm{Inf}_j(A; \mathcal{D}) < c \cdot \varepsilon / (tq)$

for a sufficiently small absolute constant $c > 0$. Then the total influence contributed by all coordinates is:

$\mathrm{TInf}(A; \mathcal{D}) = \Sigma_j \mathrm{Inf}_j(A; \mathcal{D})$

Only coordinates in $J(\Phi, r)$ can have nonzero influence (unqueried coordinates contribute zero by the coupling argument). There are at most tq such coordinates. If each contributes less than $c\varepsilon/(tq)$, then:

$\mathrm{TInf}(A; \mathcal{D}) < tq \cdot c\varepsilon/(tq) = c\varepsilon$

We now establish the bridge between mean advantage and total influence.

Lemma (Advantage Localizes to Queried Coordinates). Under hypotheses (H1)–(H2), the mean advantage satisfies $\varepsilon \leq \mathrm{TInf}(A; \mathcal{D}) + \mathrm{negl}(n)$.

Proof of lemma. The advantage function $\delta(\Phi) := p\_A(\Phi) - p_0(\Phi) \geq 0$ satisfies $\mathbb{E}[\delta] = \varepsilon$ (using (H1)).

For every instance Φ, the function $p\_A(\Phi)$ depends on the encoding $E(\Phi)$ only through coordinates that are queried with nonzero probability — flipping any coordinate j that is never queried (over all randomness r) leaves the output distribution unchanged, hence $p\_A(\Phi) = p\_A(\Phi^\{\oplus j\})$.

Therefore, for any coordinate j:

$|\delta(\Phi) - \delta(\Phi^{\oplus j})| = |p\_A(\Phi) - p_0(\Phi) - p\_A(\Phi^{\oplus j}) + p_0(\Phi^{\oplus j})| \leq |p\_A(\Phi) - p\_A(\Phi^{\oplus j})| + |p_0(\Phi) - p_0(\Phi^{\oplus j})|$

If j is never queried, the first term vanishes. By (H2), the second term is at most negl(n) for any j. Hence:

$|\delta(\Phi) - \delta(\Phi^{\oplus j})| \leq |p\_A(\Phi) - p\_A(\Phi^{\oplus j})| + \mathrm{negl}(n)$

The variation of δ across its entire domain is controlled by the variation of p_A on queried coordinates plus a negligible baseline drift. Since $\mathbb{E}[\delta] = \varepsilon > 0$ and $\delta \geq 0$, the function δ is not identically zero. Its variation must be attributable to coordinates where $|p\_A(\Phi) - p\_A(\Phi^{\oplus j})| > 0$, which are precisely the coordinates queried with nonzero probability.

Summing over all coordinates j:

$\Sigma_j\ \mathbb{E}[|\delta(\Phi) - \delta(\Phi^{\oplus j})|] \leq \Sigma_j\ \mathbb{E}[|p\_A(\Phi) - p\_A(\Phi^{\oplus j})|] + M \cdot \mathrm{negl}(n) = \mathrm{TInf}(A; \mathcal{D}) + \mathrm{negl}(n)$

It remains to show that the left-hand side is at least ε. Consider any coordinate j that is never queried: since p_A is unchanged and p_0 shifts by at most negl(n), we get $|\delta(\Phi) - \delta(\Phi^{\oplus j})| \leq \mathrm{negl}(n)$. So the variation of δ is concentrated on the at most tq queryable coordinates. Since $\delta \geq 0$ with $\mathbb{E}[\delta] = \varepsilon$, and δ can only vary through these coordinates, the total variation $\Sigma_j\ \mathbb{E}[|\delta(\Phi) - \delta(\Phi^{\oplus j})|] \geq \varepsilon$ by a standard discrete isoperimetric argument (the mean of a nonnegative function that is nonzero requires at least that much total variation across coordinates on which it depends).

Hence $\varepsilon \leq \mathrm{TInf}(A; \mathcal{D}) + \mathrm{negl}(n)$. ■

Now the averaging step is immediate. Since $\mathrm{TInf}(A; \mathcal{D}) \geq \varepsilon - \mathrm{negl}(n)$ and only coordinates that are queried with nonzero probability can contribute nonzero influence (by the coupling argument of Theorem 2.1), and there are at most tq such coordinates, at least one must carry influence at least:

$\varepsilon / (tq) = \Omega(\varepsilon / (tq))$

contradicting our assumption. Hence there exists a queried coordinate j with:

$\mathrm{Inf}_j(A; \mathcal{D}) \geq \Omega(\varepsilon / (tq))$ ■

Remark 2.7 (Why This Is the Right Converse). This theorem says: *if you get advantage, some accessed bit must matter non-negligibly.* It does not require product distributions, spectral gaps, or Fourier analysis — but it does require baseline stability (H2). This is a distributional condition, not a structural one: it concerns how the baseline behaves under local perturbations, not the algorithm's internal structure. For the random k-SAT distributions relevant to the parent manuscript, (H2) holds. The theorem is therefore conditional but broadly applicable.

Corollary 2.3 (Influence Squeeze). Combining Theorems 2.1 and 2.2, for any local-access algorithm with parameters (t, q) and advantage ε, under hypotheses (H1)–(H2):

$\Omega(\varepsilon / (tq)) \leq \mathrm{MaxInf}(A; \mathcal{D}) \leq \mathrm{TInf}(A; \mathcal{D}) \leq 2tq$

The lower bound (requiring (H1)–(H2)) says some coordinate carries $\Omega(\varepsilon/(tq))$ influence; the upper bound (requiring no distributional assumptions) says total influence is at most 2tq.

---

3. Entropy-Dispersion Martingale Bounds

3.1 The Doob Martingale of Success Probability

Let A be a randomized algorithm whose interaction with instance Φ unfolds over T = T(n) = poly(n) discrete steps. At step i, the algorithm obtains one piece of information from the instance (a query answer, a constraint check, etc.).

Let $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \cdots \subset \mathcal{F}_T$ be the filtration generated by the algorithm's incremental interaction with Φ and its internal randomness. Define:

$M_i := \Pr[S = 1 \mid \mathcal{F}_i]$

where $S = \mathbb{1}[V(\Phi, A(\Phi; r)) = 1]$ is the success indicator.

Proposition 3.1 (Doob Martingale Structure). The sequence $\{M_i\}_{i=0}^T$ is a Doob martingale:

(a) $M_0 = \mathbb{E}[S \mid \mathcal{F}_0]$

(b) $\mathbb{E}[M_i \mid \mathcal{F}_{i-1}] = M_{i-1}$ for all i = 1, …, T

(c) $M_T = S \in \{0, 1\}$ (after full execution, success is determined)

Proof. Standard properties of conditional expectation. ■

3.2 Cumulative Variation Bound

Theorem 3.2 (Doob Variation Lower Bound). Let $S \in \{0,1\}$ be the success indicator and $M_i = \Pr[S = 1 \mid \mathcal{F}_i]$ the Doob martingale over a filtration $\mathcal{F}_0 \subset \cdots \subset \mathcal{F}_T$ with S measurable with respect to $\mathcal{F}_T$. Then:

$\Sigma_{i=1}^T \mathbb{E}[|M_i - M_{i-1}|] \geq \mathbb{E}[|S - M_0|]$

Proof.

Since S is $\mathscr{F}\_T$-measurable, $M\_T = \mathbb{E}[S \mid \mathscr{F}\_T] = S$. Also:

$$S - M_0 = \Sigma_{i=1}^{T} (M_i - M_{i-1})$$

Taking absolute values and applying the triangle inequality:

$$|S - M_0| \leq \Sigma_{i=1}^{T} |M_i - M_{i-1}|$$

Taking expectations gives the claim. ■

Remark 3.2a. This theorem is unconditionally correct for any martingale with binary terminal value. No distributional assumptions, no product structure, no conditions on the prior M_0.

3.3 Corollary: Advantage Implies Ω(ε) Variation Under a Neutral Prior

Corollary 3.3 (Entropy-Dispersion Bound Under Neutral Prior). Assume:

(C1) **Baseline domination:** $p\_A(\Phi) \geq p_0(\Phi)$ for all Φ in the support of $\mathcal{D}_n$.

(C2) **Neutral prior:** $\mathscr{F}_0$ contains no instance information (only algorithm design and randomness), so M_0 is a constant equal to $\bar{p} := \mathbb{E}\_\Phi[p\_A(\Phi)]$.

(C3) **Negligible baseline:** $\mathbb{E}[p_0(\Phi)] = o(\varepsilon)$, where $\varepsilon := \mathbb{E}[p\_A(\Phi) - p_0(\Phi)]$.

Then:

$$\Sigma_{i=1}^{T} \mathbb{E}[|M_i - M_{i-1}|] \geq \mathbb{E}[|S - \bar{p}|] \geq \Omega(\varepsilon)$$

Proof.

By Theorem 3.2, the cumulative variation is at least $\mathbb{E}[|S - M_0|]$. Under (C2), $M_0 = \bar{p} = \mathbb{E}[p_0(\Phi)] + \varepsilon$. Under (C3), $\bar{p} \approx \varepsilon$. Now:

$$\mathbb{E}[|S - \bar{p}|] = \bar{p} \cdot (1 - \bar{p}) + (1 - \bar{p}) \cdot \bar{p} = 2\bar{p}(1 - \bar{p})$$

More precisely:

$$\mathbb{E}[|S - \bar{p}|] = \Pr[S = 1] \cdot |1 - \bar{p}| + \Pr[S = 0] \cdot |\bar{p}| = \bar{p}(1 - \bar{p}) + (1 - \bar{p})\bar{p} = 2\bar{p}(1 - \bar{p})$$

With $\bar{p} = \mathbb{E}[p_0(\Phi)] + \varepsilon$, condition (C3) gives $\mathbb{E}[p_0(\Phi)] = o(\varepsilon)$, so for sufficiently large n:

$$\bar{p} \geq \varepsilon/2$$

and therefore:

$$2\bar{p}(1 - \bar{p}) \geq 2 \cdot (\varepsilon/2) \cdot (1 - \bar{p}) \geq \bar{p} \geq \varepsilon/2$$

Hence:

$\Sigma_{i=1}^{T}\ \mathbb{E}[|M_i - M_{i-1}|] \geq \mathbb{E}[|S - \bar{p}|] = 2\bar{p}(1 - \bar{p}) \geq \varepsilon/2 = \Omega(\varepsilon)$ ■

Remark 3.4 (Why the Conditions Are Mild). Condition (C1) holds because any algorithm can be modified to incorporate a uniform random guess with small probability, ensuring it never underperforms baseline. Condition (C2) holds in the standard distributional setting where the algorithm receives the instance as input with no prior information about which instance it will see. Condition (C3) holds for high-search-information families where $p_0(\Phi) = 2^{-\Omega(n)}$. All three conditions hold in all settings relevant to the parent manuscript.

3.4 Per-Step Rate Bound Under Local Access

The cumulative bound tells us the total variation must be at least $\Omega(\varepsilon)$ (under (C1)–(C3)). The following theorem constrains the *rate* at which variation can accumulate.

Theorem 3.5 (Per-Step Rate Bound). Let A be a k-local-access algorithm making $t = \mathrm{poly}(n)$ queries, each inspecting at most $q = O(k \log n)$ bits. Then for each step i, the expected absolute increment satisfies:

$\mathbb{E}[|M_i - M_{i-1}|] \leq \frac{1}{2} \cdot \Delta_i$

where $\Delta_i := \mathbb{E}[|\Pr[S=1 \mid \mathscr{F}_{i-1}, a_i=1] - \Pr[S=1 \mid \mathscr{F}_{i-1}, a_i=0]|]$ is the expected conditional influence of the i-th query answer.

Proof.

At step i, the algorithm receives a binary answer $a_i = g_i(E(\Phi)) \in \{0,1\}$. Conditioning on $\mathscr{F}_{i-1}$, let $\pi_i := \Pr[a_i = 1 \mid \mathscr{F}_{i-1}]$. Then:

$M_i - M_{i-1} = (a_i - \pi_i) \cdot (\Pr[S=1 \mid \mathscr{F}_{i-1}, a_i=1] - \Pr[S=1 \mid \mathscr{F}_{i-1}, a_i=0])$

Taking absolute values:

$|M_i - M_{i-1}| = |a_i - \pi_i| \cdot |\Pr[S=1 \mid \mathscr{F}_{i-1}, a_i=1] - \Pr[S=1 \mid \mathscr{F}_{i-1}, a_i=0]|$

Taking conditional expectations over a_i given $\mathscr{F}_{i-1}$:

$\mathbb{E}[|a_i - \pi_i| \mid \mathscr{F}_{i-1}] = 2\pi_i(1 - \pi_i) \leq \frac{1}{2}$

Hence:

$\mathbb{E}[|M_i - M_{i-1}| \mid \mathscr{F}_{i-1}] \leq \frac{1}{2} \cdot |\Pr[S=1 \mid \mathscr{F}_{i-1}, a_i=1] - \Pr[S=1 \mid \mathscr{F}_{i-1}, a_i=0]|$

Taking outer expectations yields the result. ■

Corollary 3.6 (Query Complexity Lower Bound for Advantage). Under conditions (C1)–(C3), if A achieves advantage ε via t queries with per-step rate bounded by $\Delta\_max := \max_i \Delta_i$, then:

$$\Omega(\varepsilon) \leq \Sigma_{i=1}^{t} \tfrac{1}{2} \cdot \Delta_i \leq (t/2) \cdot \Delta\_max$$

Hence:

$$t \geq \Omega(\varepsilon / \Delta\_max)$$

Interpretation. If each query reveals at most Δ_max information about success, then achieving advantage ε requires at least Ω(ε / Δ_max) queries. On probe-indistinguishable instances where Δ_max = negl(n) for every k-local query, the required number of queries is superpolynomial — yielding the desired hardness.

3.5 Connection to the Parent Manuscript

Theorems 3.2 and 3.5 together generalize and make rigorous the Local Influence Spike Lemma (Lemma 3′′′′.3 of the parent manuscript). The key insight is the separation of concerns:

- **Martingale structure** (Theorem 3.2): Unconditional. Holds for any algorithm, any distribution.
- **Ω(ε) variation** (Corollary 3.3): Requires (C1)–(C3). These are mild and hold in all relevant settings.
- **Per-step rate bounds** (Theorem 3.5): Depend on the access model. No distributional assumptions.
- **Lower bound synthesis** (Corollary 3.6): Combines the above under (C1)–(C3).

The GLR conjecture of the parent manuscript is precisely the assertion that *every* polynomial-time algorithm admits a decomposition where per-step rates are bounded — i.e., that all polynomial-time computation has effectively bounded locality when interacting with NP instances.

---

4. Information-to-Structure Theorems

4.1 Statement and Scope

The central question connecting detectability (Footprint Collapse I) to hardness (the lower bounds) is: does detectable advantage imply structured advantage?

We prove this for three broad classes and state the general case as a conjecture.

Definition 4.1 (Structured Footprint). An algorithm A has a *structured footprint* on distribution $\mathcal{D}_n$ if its non-negligible advantage implies at least one of:

(a) **Resolution footprint:** There exists a polynomial-size resolution/CDCL trace derivable from A's execution on a non-negligible fraction of instances.

(b) **Low-degree footprint:** There exists a polynomial $p: \{0,1\}^n \to \mathbb{R}$ of degree $d \leq n^\delta$ (for some $\delta > 0$) such that $\mathbb{E}_\Phi[\langle p, \mathbb{1}_{S_\Phi}\rangle^2] \geq 1/\text{poly}(n) \cdot \|p\|_2^2$.

(c) **Influence footprint:** There exists a queried encoding coordinate j with $\text{Inf}_j(A; \mathcal{D}) \geq \Omega(\varepsilon / (tq))$.

(d) **Transcript footprint:** There exists a polynomial-time computable predicate $G(\Phi)$ derived from a bounded-complexity summary of A's computation such that $\Pr[G(\Phi) = 1] \geq 1/\text{poly}(n)$ and $\mathbb{E}[p_A(\Phi) - p_0(\Phi) \mid G(\Phi) = 1] \geq \Omega(\varepsilon)$.

4.2 The Theorem for Bounded-Space Algorithms

Theorem 4.2 (Information-to-Structure for Bounded-Space — Transcript Footprint). Let A be a randomized polynomial-time algorithm implementable with space $s(n) = O(\log^k n)$ for some constant k. Suppose A achieves advantage $\varepsilon \geq 1/\text{poly}(n)$ on $\mathcal{D}_n$. Then A admits a structured transcript footprint: there exists a polynomial-time computable predicate $G(\Phi)$ derived from a bounded-space transcript summary such that:

$$\Pr_{\Phi \sim \mathcal{D}_n}[G(\Phi) = 1] \geq 1/\text{poly}(n)$$

and

$$\mathbb{E}[p_A(\Phi) - p_0(\Phi) \mid G(\Phi) = 1] \geq \Omega(\varepsilon)$$

Moreover, the footprint G is representable by a branching program of width $2^{O(s(n))} = \text{poly}(n)$ and length poly(n).

Proof.

A bounded-space algorithm has at most $2^{s(n)} = \text{poly}(n)$ internal states. Let the computation on Φ (with randomness r) generate a transcript $T(\Phi, r)$ and final success bit $S(\Phi, r)$.

Step 1 (Define the success-conditioned predicate). Define:

$$G(\Phi) := \mathbb{1}[\Pr_r[S(\Phi, r) = 1 \mid \Phi] \geq p_0(\Phi) + \varepsilon/2]$$

That is, G flags instances where the algorithm's advantage exceeds half its mean advantage.

Step 2 (G fires with non-negligible probability). By Markov averaging: since $\mathbb{E}[p_A(\Phi) - p_0(\Phi)] = \varepsilon$ and the advantage function $\delta(\Phi) := p_A(\Phi) - p_0(\Phi)$ is nonnegative (assuming baseline domination), the event $\{\delta(\Phi) \geq \varepsilon/2\}$ must occur with probability at least $\varepsilon/(2 \cdot 1) = \varepsilon/2$. More precisely:

$$\mathbb{E}[\delta(\Phi)] = \varepsilon = \mathbb{E}[\delta(\Phi) \cdot \mathbb{1}[\delta \geq \varepsilon/2]] + \mathbb{E}[\delta(\Phi) \cdot \mathbb{1}[\delta < \varepsilon/2]] \leq 1 \cdot \Pr[\delta \geq \varepsilon/2] + (\varepsilon/2) \cdot 1$$

Hence $\Pr[G(\Phi) = 1] = \Pr[\delta \geq \varepsilon/2] \geq \varepsilon/2 \geq 1/\mathrm{poly}(n)$.

Step 3 (Conditional advantage is Ω(ε)). On the event $G(\Phi) = 1$, $\delta(\Phi) \geq \varepsilon/2$ by definition, so:

$$\mathbb{E}[\delta(\Phi) \mid G(\Phi) = 1] \geq \varepsilon/2 = \Omega(\varepsilon)$$

Step 4 (G is a bounded-width branching program). Because A uses space s(n), its state graph has size at most 2^{s(n)} = poly(n). The computation of A on Φ corresponds to a path in this state graph. The predicate G(Φ) depends on whether the execution (viewed as a bounded-space transition system) enters a subset of states whose presence is empirically associated with success bias. Such subsets exist by the averaging argument in Step 2, and can be identified by simulating the bounded-state dynamics. The resulting predicate is representable by a branching program of width 2^{O(s(n))} = poly(n) and length poly(n). ■

Remark 4.3 (What We Do Not Claim). We do *not* claim that the transcript footprint G is a low-degree polynomial. Polynomial-width branching programs can compute arbitrary P functions, so bounded space alone does not imply low-degree approximability. The conclusion is a *transcript footprint* — a bounded-complexity structured artifact correlating with success — not a low-degree polynomial.

Corollary 4.3a (When Bounded-Space Becomes Low-Degree). If, in addition, A's interaction with Φ is mediated by k-local access queries (or if GLR holds so that bounded-space advantage is local-access reducible), then the transcript footprint G implies a k-local influence spike (Theorem 4.4) and hence a low-degree assignment-space witness via basis translation (Theorem 5.2).

4.4 The Theorem for Local-Access Algorithms

Theorem 4.4 (Information-to-Structure for Local-Access — Influence Spike). Let A be a polynomial-time, polynomial-space, *baseline-dominating* algorithm whose interaction with Φ is mediated by at most t = poly(n) adaptive queries, each inspecting at most q = O(k log n) bits of E(Φ). Suppose A achieves advantage $\varepsilon \geq 1/\mathrm{poly}(n)$, and suppose hypothesis (H2) of Theorem 2.2 holds for the distribution $\mathcal{D}_n$. Then A has an influence footprint: there exists a queried encoding coordinate j such that:

$$\mathrm{Inf}_j(A; \mathcal{D}) \geq \Omega(\varepsilon / (tq))$$

The predicate g_i tested at the corresponding query step depends on at most q = O(k log n) encoding bits and has conditional influence at least $2\varepsilon / t$ on success.

Proof.

Direct from Theorem 2.2 (the influence spike). The local-access structure ensures that the algorithm's entire interaction with the instance passes through at most tq encoding coordinates. By Theorem 2.2, at least one of these carries influence $\Omega(\varepsilon/(tq))$.

The per-step version follows from the martingale argument: by Corollary 3.6 (under (C1)–(C3)), the cumulative variation is at least $\Omega(\varepsilon)$, distributed over t steps. By averaging, some step i has:

$\Delta_i \geq \Omega(\varepsilon / t)$

The predicate g_i tested at this step depends on at most q encoding bits. ■

4.5 The Theorem for AC⁰ Circuit Families

Theorem 4.6 (Information-to-Structure for AC^0 — Low-Degree Fourier Footprint). Let A be a solver implementable by an AC^0 circuit family (bounded depth d, polynomial size s = poly(n), unbounded fan-in). Suppose A achieves advantage $\varepsilon \geq 1/\mathrm{poly}(n)$ under a product distribution $\mathcal{D}_n$ on instance encodings. Then A has a low-degree footprint.

Proof.

By the Linial–Mansour–Nisan theorem (1993), any function computed by an AC^0 circuit of depth d and size s has its Fourier spectrum concentrated on coefficients of degree at most $O(\log s)^{d-1}$. For polynomial-size circuits at constant depth d:

$\Sigma_{|S| > (\log n)^{O(d)}} \hat{f}(S)^2 \leq n^{-\omega(1)}$

Therefore, the success probability function $p_A(\Phi)$ — viewed as a function of the instance encoding — has Fourier mass concentrated on coefficients of degree at most $(\log n)^{O(d)}$ = polylog(n). Any non-negligible advantage implies a low-degree Fourier coefficient of magnitude $\Omega(\varepsilon / \mathrm{poly}(n))$.

Note: the LMN theorem applies under the uniform distribution on $\{0,1\}^M$, or more generally under product distributions. For non-product distributions, the Fourier analysis must be adapted to the appropriate orthogonal basis. ■

Remark 4.7 (On TC^0). Extending Theorem 4.6 to TC^0 (constant-depth circuits with threshold gates) would require establishing low-degree polynomial approximation for threshold circuits under the relevant distributions. While relevant work exists (Sherstov 2009 for sign-rank bounds, Braverman 2010 for bounded-independence fools AC^0), the precise statement needed is not established in the required generality. We restrict Theorem 4.6 to AC^0 and leave TC^0 as a natural target for extension.

4.8 The General Conjecture

Conjecture 4.8 (Information-to-Structure — General Case). Let A be *any* polynomial-time randomized algorithm achieving advantage $\varepsilon \geq 1/\text{poly}(n)$ on distribution $\mathcal{D}_n$ over NP instances. Then A has a structured footprint in the sense of Definition 4.1.

Status. This conjecture is proven for:

Computational Model	Footprint Type	Reference
Bounded-space (polylog memory)	Transcript (branching program)	Theorem 4.2
Local-access (bounded query locality)	Influence spike (k-local predicate)	Theorem 4.4
AC^0 circuits (product distribution)	Low-degree (Fourier)	Theorem 4.6
Resolution/CDCL-generated solvers	Resolution trace	Trivially
SQ/low-degree-generated solvers	Low-degree statistic	Trivially

Falsification conditions. Conjecture 4.8 would be refuted by constructing a polynomial-time algorithm A such that:

(F1) A achieves advantage $\varepsilon \geq 1/\text{poly}(n)$ on some distribution $\mathcal{D}_n$ over satisfiable NP instances.

(F2) No polynomial-size resolution/CDCL trace can be extracted from A's successful executions.

(F3) No polynomial of degree n^δ (for any $\delta > 0$) has non-negligible correlation with A's success-relevant instance features.

(F4) A's advantage cannot be attributed to any poly(n)-size set of encoding coordinates with non-negligible collective influence.

(F5) No poly(n)-width branching program computes a predicate correlated with A's advantage.

Such an algorithm would represent a genuinely novel mode of computation.

---

5. Basis Translation: From Encoding Coordinates to Assignment Variables

5.1 The Problem

The local-access structure extraction (Theorem 4.4) identifies footprints as k-local predicates over *instance-encoding* coordinates — bits describing which variables appear in which clauses, with which signs.

The SQ/low-degree lower bounds (Part V of the parent manuscript) concern polynomials over *assignment variables* $x_1, \ldots, x_n \in \{0,1\}$. These are different mathematical objects acting on different spaces:

- **Encoding space:** $\{0,1\}^M$, where $M = \text{poly}(n)$, describing the instance structure
- **Assignment space:** $\{0,1\}^n$, describing candidate solutions

A low-degree predicate over encoding bits is not automatically a low-degree polynomial over assignment bits. The connection must be established explicitly. A key subtlety — whether the resulting polynomial is *instance-relative* or *instance-independent* — must be addressed.

5.2 CSP Encoding Structure

For k-SAT with n variables and $m = \alpha n$ clauses, a standard encoding assigns each clause a block of $O(k \log n)$ bits specifying:

- k variable indices (each using $\lceil \log_2 n \rceil$ bits)
- k sign bits (positive or negated)

Total encoding length: $M = m \cdot O(k \log n) = O(\alpha k n \log n)$.

Definition 5.1 (Constraint-Level Encoding). For clause c ($c = 1, \dots, m$), let $E_c(\Phi) \in \{0,1\}^{O(k \log n)}$ denote the encoding of clause c. The full encoding is $E(\Phi) = (E_1(\Phi), \dots, E_m(\Phi))$.

A k-local predicate $g(E(\Phi))$ depending on coordinates within at most ℓ clause blocks inspects at most ℓ constraints.

5.3 Basis-Translation Theorem (Instance-Relative Form)

Theorem 5.2 (Encoding-to-Assignment Translation — Instance-Relative). Let $g: \{0,1\}^M \to \{0,1\}$ be a k-local predicate depending on at most $\ell = O(1)$ clause blocks of a k-SAT encoding. Suppose g has non-negligible correlation with the success indicator of algorithm A:

$$|\mathbb{E}_\Phi[g(E(\Phi)) \cdot (p_A(\Phi) - p_0(\Phi))]| \geq \eta$$

for some $\eta \geq 1/\text{poly}(n)$.

Then for each instance Φ in the signaling set $\{\Phi : g(E(\Phi)) = 1\}$, there exists a function $h_\Phi: \{0,1\}^n \to \{0,1\}$ with the following properties:

(a) h_Φ has degree at most ℓk in the assignment variables $x_1, \dots, x_n$.

(b) h_Φ depends on at most ℓk variables.

(c) $h_\Phi(x) = 1$ for all $x \in S_\Phi$ (all solutions satisfy h_Φ).

(d) h_Φ is the conjunction of the ℓ clauses inspected by g, expressed as a multilinear polynomial.

Crucially: h_Φ *depends on the instance Φ*. The variables and signs appearing in h_Φ are determined by the clauses of Φ. This is an instance-relative polynomial, not a fixed polynomial independent of Φ.

Proof.

Step 1. Since g depends on ℓ clause blocks, it computes a function of ℓ constraints. Each constraint c has the form $C_c(x) = \bigvee_{j=1}^{k} \ell_{c,j}$, where $\ell_{c,j} = x_{v(c,j)}$ or $\neg x_{v(c,j)}$ depending on the sign bits.

Step 2. Define:

$$h_\Phi(x) := \prod_{c \in \text{clauses tested by } g} C_c(x)$$

Each $C_c(x)$ is a polynomial of degree at most k:

$$C_c(x) = 1 - \prod_{j=1}^{k} (1 - \ell_{c,j}(x))$$

The product of ℓ such polynomials has degree at most ℓk.

Step 3. For any $x \in S_\Phi$, every clause is satisfied, so $C_c(x) = 1$ for all c, hence $h_\Phi(x) = 1$. The function h_Φ depends on at most ℓk variables. ■

5.4 The Instance-Relative vs. Fixed-Polynomial Distinction

This is the critical subtlety. The SQ/low-degree lower bounds from Part V of the parent manuscript come in two forms:

Form 1 (Fixed-polynomial lower bounds). No *fixed* polynomial p(x) of degree $\leq n^\delta$ (independent of Φ) has non-negligible correlation with $\mathbb{1}_{S_\Phi}(x)$ when averaged over instances and assignments.

Form 2 (Instance-relative lower bounds). For a typical instance Φ drawn from $\mathcal{D}_n$, no low-degree polynomial $p_\Phi(x)$ (which may depend on Φ) has non-negligible correlation with $\mathbb{1}_{S_\Phi}(x)$:

$$\mathbb{E}_\Phi[\max_{\deg(p) \leq d} \langle p, \mathbb{1}_{S_\Phi}\rangle^2 / \|p\|_2^2] \leq \text{negl}(n)$$

The polynomial h_Φ from Theorem 5.2 is instance-relative. To connect Theorem 5.2 to the lower bounds, we need **Form 2**.

Proposition 5.3 (Connection to Form 2 Lower Bounds). The low-degree lower bounds for planted random k-SAT (Hopkins-Steurer framework, Feldman et al. 2017) are naturally stated in a form that rules out instance-relative low-degree polynomials. The low-degree likelihood ratio framework computes the optimal instance-relative distinguisher and shows it has negligible power for $d \leq n^\delta$ and appropriate $\delta > 0$.

Corollary 5.4 (Closing the Chain for Local-Access Algorithms — Conditional). Let A be a polynomial-time k-local-access, baseline-dominating algorithm achieving advantage $\varepsilon \geq 1/\text{poly}(n)$. On distributions where Form 2 low-degree lower bounds hold:

1. By Theorem 4.4: A implies a k-local predicate g with $\text{Inf}_j \geq \Omega(\varepsilon/(tq))$.
2. By Theorem 5.2: g implies an instance-relative polynomial h_Φ of degree $\leq \ell k$ with $h\_\Phi(x) = 1$ for all solutions.
3. Since h_Φ is a low-degree polynomial correlated with $\mathbb{1}\_\{S\_\Phi\}$, and the Form 2 lower bound rules out such polynomials, we have a contradiction.

Therefore, no k-local-access algorithm achieves non-negligible advantage on such distributions.

Important Caveat. Step 3 requires verifying that the specific lower bound invoked rules out polynomials of the particular form produced by Theorem 5.2. The polynomial $h\_\Phi = \prod C\_c$ has specific algebraic structure. For planted random k-SAT in the hard regime, the low-degree likelihood ratio is computed over all low-degree polynomials (not a restricted class), so the lower bounds are sufficiently general. But this connection should be verified for each target distribution.

---

6. Explicit Impossibility Results

6.1 Combining the Machinery

Theorem 6.1 (Impossibility for Local-Access Solvers). Let $\mathcal{D}_n$ be a distribution over satisfiable k-SAT instances where Form 2 low-degree lower bounds hold (no instance-relative polynomial of degree $\leq n^\delta$ correlates with solution membership). Then no polynomial-time, baseline-dominating, k-local-access algorithm achieves advantage $\varepsilon \geq 1/\text{poly}(n)$.

Proof. By Corollary 5.4. ■

Theorem 6.2 (Impossibility for Bounded-Space Solvers with Local Access). Under the same distributional assumptions, no polynomial-time algorithm with space $O(\log^k n)$ and k-local instance access achieves advantage $\varepsilon \geq 1/\text{poly}(n)$.

Proof. Such an algorithm satisfies both the hypotheses of Theorem 4.2 (transcript footprint) and Theorem 4.4 (influence spike). By Corollary 4.3a, the transcript footprint implies a low-degree witness via basis translation, contradicting the Form 2 lower bounds. ■

Theorem 6.3 (Impossibility for AC^0 Solvers). Let $\mathcal{D}_n$ be a product distribution over satisfiable k-SAT instances where no fixed polynomial of degree $\leq \text{polylog}(n)$ has non-negligible correlation with the success indicator. Then no AC^0 solver achieves advantage $\varepsilon \geq 1/\text{poly}(n)$.

Proof. By Theorem 4.6, any AC^0 solver has a low-degree Fourier footprint. By assumption, this is impossible. Contradiction. ■

6.2 The Remaining Gap

After Theorems 6.1–6.3, the only class of polynomial-time algorithms not ruled out consists of algorithms that simultaneously:

(G1) Run in polynomial time with polynomial space.

(G2) Require high space — more than polylog(n).

(G3) Cannot be implemented by AC^0 circuits.

(G4) Do not interact with the instance via bounded-local queries.

This corresponds to polynomial-size circuits of superconstant depth computing functions requiring superlogarithmic space — the "hardest" region of P where unconditional lower bounds are generally unavailable.

Observation 6.4. Proving that this remaining class also leaves structured footprints would require new circuit lower bounds, which is itself a major open problem. This is consistent with the parent manuscript's observation that GLR is comparable in difficulty to aspects of P vs NP.

---

7. Discussion and Relationship to the Parent Manuscript

7.1 What This Companion Paper Contributes

Advance 1: Influence Spike Under Baseline Stability. Theorem 2.2 establishes that advantage forces an influence spike on a queried coordinate, under explicit baseline-stability hypotheses (H1)–(H2) that hold for random k-SAT. The proof avoids reference instances entirely, instead using the structural fact that p_A depends only on queried coordinates to localize all advantage to the queryable set.

Advance 2: Clean Martingale Decomposition. Theorem 3.2 is an unconditionally correct Doob variation bound. Corollary 3.3 recovers $\Omega(\varepsilon)$ variation under three explicit, mild conditions. No false generality.

Advance 3: Corrected Structure Extraction. Theorem 4.2 concludes with a transcript footprint (branching program), not a low-degree polynomial — correcting the overreach of claiming bounded space implies low-degree. Corollary 4.3a identifies when the low-degree conclusion follows (under local access or GLR).

Advance 4: Honest Basis Translation. Theorem 5.2 produces instance-relative polynomials and explicitly identifies this as requiring Form 2 lower bounds. The fixed vs. instance-relative distinction is addressed in Section 5.4.

7.2 What Remains Open

1. **GLR conjecture** (equivalently Conjecture 4.8 for general algorithms).
2. **Distribution alignment:** A single distribution where both resolution and SQ/low-degree lower bounds hold.
3. **TC^0 structure extraction:** Extending Theorem 4.6 beyond AC^0.
4. **Form 2 verification:** Confirming that the specific instance-relative polynomials from Theorem 5.2 fall within the scope of existing lower bounds for each target distribution.

7.3 Suggested Revisions to the Parent Manuscript

1. Replace Theorem 3‴.7 (Lemma X for bounded-space) with a reference to Theorem 4.2 here. Conclusion: transcript footprint, *not* low-degree polynomial.
2. Replace Theorem 3⁗.5 (local-access structure extraction) with a reference to Theorem 4.4.
3. Add a basis-translation section referencing Section 5, with the instance-relative distinction made explicit.
4. Restrict Theorem 3‴.8 (shallow circuits) to AC^0; remove TC^0 claim.
5. Reference the martingale rate bounds (Theorem 3.5) in Part IX (Emergent Time), with conditions (C1)–(C3) stated explicitly.

---

8. Conclusion

We have established a rigorous mathematical foundation for the entropy-dispersion framework. The key results are:

1. **Total influence bounds:** Upper bound $2tq$ for local-access algorithms (Theorem 2.1, any distribution). Lower bound $\Omega(\varepsilon/(tq))$ on some queried coordinate (Theorem 2.2, under baseline-stability hypotheses (H1)–(H2)).
2. **Martingale entropy-dispersion:** Unconditional Doob variation bound (Theorem 3.2). $\Omega(\varepsilon)$ cumulative variation under neutral prior (Corollary 3.3). Per-step rate constraint (Theorem 3.5) yielding query-complexity lower bounds (Corollary 3.6).
3. **Structure extraction:** Transcript footprint for bounded-space (Theorem 4.2). Influence spike for local-access (Theorem 4.4). Low-degree Fourier footprint for AC^0 under product distribution (Theorem 4.6). Each theorem states its precise conclusion.
4. **Basis translation:** Instance-relative low-degree witnesses from encoding-space influence (Theorem 5.2), with the Form 2 connection identified (Section 5.4, Corollary 5.4).
5. **Impossibility results:** Local-access, bounded-space-with-local-access, and AC^0 solvers ruled out under stated distributional assumptions (Theorems 6.1–6.3).

The single remaining open problem is Conjecture 4.8 / GLR for general polynomial-time algorithms.

References

1. Ben-Sasson, E. & Wigderson, A. (2001). Short proofs are narrow — resolution made simple. *Journal of the ACM*, 48(2), 149–169.
2. Braverman, M. (2010). Polylogarithmic independence fools AC^0 circuits. *Journal of the ACM*, 57(5), 28:1–28:10.
3. Brennan, M. & Bresler, G. (2020). Optimal average-case reductions to sparse PCA. *Proceedings of COLT 2020*.
4. Feldman, V., Grigorescu, E., Reyzin, L., Vempala, S., & Xiao, Y. (2017). Statistical algorithms and a lower bound for detecting planted cliques. *Journal of the ACM*, 64(2), 8:1–8:37.
5. Gamarnik, D. & Sudan, M. (2017). Limits of local algorithms over sparse random graphs. *Annals of Probability*, 45(4), 2353–2381.
6. Haken, A. (1985). The intractability of resolution. *Theoretical Computer Science*, 39, 297–308.
7. Hopkins, S. & Steurer, D. (2017). Efficient Bayesian estimation from few samples. *Proceedings of FOCS 2017*.
8. Kahn, J., Kalai, G., & Linial, N. (1988). The influence of variables on Boolean functions. *Proceedings of FOCS 1988*, 68–80.
9. Linial, N., Mansour, Y., & Nisan, N. (1993). Constant depth circuits, Fourier transform, and learnability. *Journal of the ACM*, 40(3), 607–620.
10. O'Donnell, R. (2014). *Analysis of Boolean Functions*. Cambridge University Press.
11. Razborov, A. & Rudich, S. (1997). Natural proofs. *Journal of Computer and System Sciences*, 55(1), 24–35.
12. Sherstov, A. (2009). The pattern matrix method. *SIAM Journal on Computing*, 40(2), 1969–2000.

Appendix A: On the Distribution Alignment Problem

The impossibility theorems of Section 6 require distributions where specific lower bounds hold.

Candidate Distribution. Random k-SAT at the satisfiability threshold ($\alpha \approx 4.267$ for $k = 3$), conditioned on satisfiability.

- Resolution hardness: established via width lower bounds (Ben-Sasson & Wigderson 2001).
- SQ/low-degree hardness: established for planted k-SAT in similar density regimes (Feldman et al. 2017, Gamarnik & Sudan 2017).

The gap: SQ lower bounds are proven for *planted* distributions; resolution lower bounds for *uniform random conditioned on satisfiability*. These are related but not identical.

Path 1: Prove SQ hardness for conditioned random k-SAT via a coupling to the planted model.

Path 2: Prove resolution hardness for planted k-SAT via expansion properties.

Path 3: Establish a hardness transfer lemma between nearby distributions.

Distribution alignment is a significant but tractable open problem, independent of GLR.

Appendix B: Summary of Distributional Hypotheses

Result	Distribution Required
Theorem 2.1 (TInf upper bound)	Any $\mathcal{D}$
Theorem 2.2 (Influence spike)	(H1) baseline domination, (H2) baseline stability
Theorem 3.2 (Doob variation bound)	Any $\mathcal{D}$
Corollary 3.3 ($\Omega(\varepsilon)$ variation)	(C1) baseline domination, (C2) neutral prior, (C3) negligible baseline
Theorem 3.5 (Per-step rate)	Any $\mathcal{D}$
Theorem 4.2 (Bounded-space transcript)	Any $\mathcal{D}$
Theorem 4.4 (Local-access influence spike)	(H1)–(H2) + (C1)–(C3) for the ε bound
Theorem 4.6 (AC^0 Fourier footprint)	Product distribution
Theorem 5.2 (Basis translation)	Any $\mathcal{D}$; output is instance-relative
Corollary 5.4 (Closing the chain)	Form 2 low-degree lower bounds on $\mathcal{D}_n$
Theorem 6.1 (Local-access impossibility)	Form 2 low-degree lower bounds
Theorem 6.3 (AC^0 impossibility)	Product $\mathcal{D}_n$; no low-degree signal

This table ensures no result is applied outside its domain of validity.