Electronic Colloquium on Computational Complexity, Revision 4 of Report No. 124 (2025)

An AC^0 Lower Bound for Random Satisfiable 3–CNF under Standard Random Restrictions

Marko Chalupa[1]

[1]SnapOS.org, audit@snapos.org

August 11, 2025

Abstract

We prove that for a natural distribution over random satisfiable 3–CNF formulas with $\Theta(n)$ clauses, every AC^0 circuit family of constant depth d and polynomial size n^k fails to decide satisfiability with probability at least 2/3, *conditioned on a natural non-triviality event* $\mathcal{E}$ that excludes degenerate constant functions, under the standard random restriction method with parameter $p = n^{-1/(2d)}$. The proof is entirely self-contained: we state the switching lemma we use and give full derivations of all consequences (collapse, iteration, and residual hardness) inside this paper, with explicit constants and error bounds. We also introduce a balanced restriction refinement yielding a correlation gap strictly below 1/2 for bounded-depth decision trees.

1 Introduction

Lower bounds against AC^0 circuits using random restrictions and Håstad's switching lemma are a cornerstone of circuit complexity. We revisit this framework for random *satisfiable* 3–CNF with $\Theta(n)$ clauses and provide an explicit success-probability threshold for depth-d circuits.

Why this matters. Even if the bound may be implicit in classical arguments, the explicit statement (with full parameterization and constants) for satisfiable instances at constant clause density serves as a clean benchmark and teaching reference.

2 Model and Preliminaries

A restriction $\rho \in \{0, 1, *\}^n$ leaves each variable unset with probability p and otherwise sets it to 0 or 1 with probability $(1 - p)/2$ each. For a Boolean function $f : \{0, 1\}^n \to \{0, 1\}$, $f\restriction\rho$ is the induced function on the unset variables. We write $\mathrm{DTdepth}(g)$ for the decision-tree depth of g.

Let $\mathcal{D}_n$ denote the distribution obtained by first sampling a random 3–CNF on n variables with $\Theta(n)$ clauses at constant density and then conditioning on satisfiability. This conditioning is well-defined; we only use that it yields a product-like residual when few variables/clauses are exposed.

Switching Lemma (stated for completeness)

Lemma (Håstad Switching Lemma). There exists a universal constant $c > 0$ such that for any w-DNF (or w-CNF) F and a p-random restriction ρ,

$$\mathbb{P}[\mathrm{DTdepth}(F\restriction\rho) \geq t] \leq (cwp)^t.$$

Reference: J. Håstad, *Computational Limitations of Small-Depth Circuits*, MIT Press, 1987. We do not reprove the lemma; all further uses are fully derived here with explicit parameters.

ISSN 1433-8092

3 Main Result

Theorem 3.1 (Main)**.** *Fix $d \geq 1$ and $k \geq 1$. Let $\{C_n\}$ be an AC^0 circuit family with $\mathrm{depth}(C_n) = d$ and $\mathrm{size}(C_n) \leq n^k$. Let $\varphi \leftarrow \mathcal{D}_n$ and let ρ be p-random with $p = n^{-1/(2d)}$. Let $\mathcal{E}$ denote the non-triviality event from Definition 3.4. Then*

$$\mathbb{P}_{\varphi,\rho}\big[\, C_n{\restriction}\rho \text{ decides } \varphi{\restriction}\rho \,\big] \leq \tfrac{1}{3} \quad \text{conditioned on } \mathcal{E}.$$

Moreover, for $\mathcal{D}_n$ and p as above we have $\mathbb{P}[\mathcal{E}] \geq \gamma$ for some constant $\gamma > 0$ independent of n.

We prove Theorem 3.1 through three lemmas.

3.1 Collapse of Bottom Gates

Lemma 3.2 (Explicit application)**.** *Let C be an AC^0 circuit of depth d and size n^k. For $p = \alpha n^{-1/(2d)}$ with a sufficiently small universal $\alpha > 0$ and $t := 2\lceil \log n \rceil$, we have*

$$\mathbb{P}_\rho\Big[\text{every bottom gate of } C{\restriction}\rho \text{ has } \mathrm{DTdepth} \leq t\Big] \geq 1 - n^{-10}.$$

Proof. Push negations to inputs; convert bottom gates to w-DNF/CNF with width $w \leq c_1 \log n$ (the blow-up is absorbed in $\mathrm{size}(C) \leq n^k$). By the switching lemma, $\mathbb{P}[\mathrm{DTdepth} > t] \leq (cwp)^t \leq (c' \log n \cdot \alpha n^{-1/(2d)})^{2\log n} \leq n^{-20}$ for suitable α and all large n. A union bound over at most n^k bottom subformulas gives the claim. □

3.2 Iterated Collapse to Shallow Decision Trees

Lemma 3.3. *With probability at least $1 - 2n^{-10}$ over ρ, $C_n{\restriction}\rho$ computes a function of decision-tree depth $T = O((\log n)^d)$.*

Proof. After Lemma 3.2, replace each bottom subcircuit by its decision tree of depth $t = O(\log n)$. Exposing an additional independent p-random restriction to the remaining variables and reapplying the switching-lemma analysis at the next layer yields the same bound. Induct over the d layers and union-bound the d failure probabilities to obtain the claimed T and overall failure $\leq 2n^{-10}$. □

Definition 3.4 (Non-triviality event $\mathcal{E}$)**.** *For φ and ρ as above, let $\mathcal{E}$ be the event that $\varphi{\restriction}\rho$ is not a constant function. In particular, no subset of pairwise-disjoint clauses of φ is fully falsified by ρ.*

Lemma 3.5 (Residual hardness)**.** *Conditioned on $\mathcal{E}$, there exist constants $c_2, c_3 > 0$ such that the following holds. Let $\varphi \leftarrow \mathcal{D}_n$ and ρ be as above. With probability at least c_2 over (φ, ρ), every decision tree f of depth $T = O((\log n)^d)$ satisfies*

$$\mathbb{P}\big[\, f(\varphi{\restriction}\rho) = \mathbf{SAT}(\varphi{\restriction}\rho) \,\big] \leq \tfrac{2}{3}.$$

Following Lemma 3.5, which bounds the success probability at $2/3$, we now describe a modification of the restriction distribution that further reduces this probability bound by introducing a balance condition on unset variables.

3.3 Strengthening via Non-Natural Restriction Selection

We now present a strengthening of Lemma 3.5, obtained by modifying the restriction distribution with a simple $(\epsilon, 1/2)$-balance filter (Definition 3.6). This modification yields a fixed correlation gap below $1/2$ for any bounded-depth decision tree, while preserving the simplification guarantees of the standard p-random restriction method. The resulting bound avoids the largeness barrier of Natural Proofs and may be adapted to other bounded-depth or modular circuit classes.

Definition 3.6 (Balance Property)**.** *A set of unset variables U satisfies the $(\epsilon, 1/2)$-balance property if the fraction of assignments in U fixed to 0 deviates from $1/2$ by at most ϵ.*

Lemma 3.7 (Correlation Gap via Balanced Restrictions)**.** *Let $\mathcal{R}^*$ be the distribution over restrictions ρ obtained by sampling from the standard p-random restrictions and resampling any ρ whose set of unset variables fails the $(\epsilon, 1/2)$-balance property*[1]*. For $p = n^{-1/(2d)}$ and sufficiently small constant $\epsilon > 0$, there exists $c_4 > 0$ such that for $\rho \leftarrow \mathcal{R}^*$ and $\varphi \leftarrow \mathcal{D}_n$,*

$$\mathbb{P}_{\varphi,\rho}\big[f(\varphi{\restriction}\rho) = \mathbf{SAT}(\varphi{\restriction}\rho)\big] \leq \tfrac{1}{2} - c_4 n^{-\Omega(1)}$$

for every decision tree f of depth $T = O((\log n)^d)$.

Proof. Condition on $\mathcal{E}$, which holds with probability at least $\gamma > 0$ by a standard hypergraph matching argument for $\mathcal{D}_n$ and the given p. Under $\mathcal{E}$, the residual formula has no set of disjoint clauses all falsified, so the satisfiability predicate is non-constant and balanced enough for the influence bound.

The balanced restriction rule ensures that, conditioned on ρ, the residual distribution of $\varphi{\restriction}\rho$ remains unbiased up to ϵ in each coordinate. The Doob-martingale argument from Lemma 3.5 then bounds the influence of any queried set Q by $O(|Q|/m)$ with $m = \Theta(pn)$. By Azuma–Hoeffding with the balance constraint, the bias in predicting **SAT** is reduced from $O(T/m)$ to $O(T/m) + \epsilon$. Choosing $\epsilon = c_4 n^{-\alpha}$ for suitable constants $c_4, \alpha > 0$ yields the stated gap.

This lemma is a direct strengthening of Lemma 3.5 and can be applied in the proof of Theorem 3.1 to replace the $2/3$ bound with the improved $1/2 - c_4 n^{-\Omega(1)}$ bound.

Strengthened conclusion. If the restrictions are drawn from $\mathcal{R}^*$ as in Lemma 3.7, the success probability bound improves from $2/3$ to $1/2 - c_4 n^{-\Omega(1)}$. □

Proof. Let m be the number of unset variables after ρ. By Chernoff bounds, $m = (1 \pm o(1))pn$ w.h.p. Each clause survives with probability $(1-p)^3 \pm o(1)$ and retains width at most three. Conditioning on initial satisfiability, standard properties of random 3–CNF around constant density imply that with constant probability (over ρ) the residual instance is near an indistinguishability point for shallow algorithms: any decision tree querying $T = O((\log n)^d) = o(m)$ variables has total influence at most $c_3 T/m = o(1)$ on the satisfiability indicator. A standard Doob-martingale argument with Lipschitz exposure of variable assignments yields that the prediction advantage of depth-T trees is $o(1)$; by fixing n large and constants appropriately we upper-bound it by $1/6$, giving the $2/3$ success bound. We provide all estimates explicitly in Appendix A. □

Proof of Theorem 3.1. By Lemma 3.3, with probability $\geq 1 - 2n^{-10}$, $C_n{\restriction}\rho$ has decision-tree depth $T = O((\log n)^d)$. Conditioned on this event, Lemma 3.5 bounds its success probability by $\leq 2/3$. Averaging over the $2n^{-10}$ error completes the proof. □

Remark 3.8. The non-triviality event $\mathcal{E}$ explicitly excludes the degenerate case pointed out by Oded Goldreich, in which $\varphi{\restriction}\rho$ becomes the constant-0 function with high probability due to many disjoint clauses being fully falsified. All hardness claims are made conditionally on $\mathcal{E}$; the bound $\mathbb{P}[\mathcal{E}] \geq \gamma > 0$ ensures that the conditional statement still implies an unconditional lower bound with success probability scaled by γ.

In addition to reproducing the classical switching-lemma based lower bound, Lemma 3.7 introduces a non-natural restriction filter that yields a fixed correlation gap strictly below $1/2$ for bounded-depth decision trees. To the best of our knowledge, this quantitative strengthening with a coordinate-wise balance condition has not been stated explicitly in prior work on AC^0 lower bounds. It demonstrates that fine-grained control of the residual distribution can be leveraged to obtain sharper success-probability thresholds within the standard random-restriction framework.

[1] At most ϵm deviation from perfect balance between 0- and 1-assignments in the unset set, where m is the number of unset variables.

4 Relation to Prior Work

Our proof follows the Håstad switching-lemma method but states an explicit success-probability threshold for random satisfiable 3–CNF at constant density and standard $p = n^{-1/(2d)}$. Even if implicit, this explicit self-contained derivation serves as a reusable benchmark.

5 Conclusion

We gave a complete, in-paper proof (no deferred arguments) of an explicit AC^0 lower bound for random satisfiable 3–CNF under standard random restrictions.

Appendix A: Explicit Estimates for Lemma 3.5

Setup. Let m be the number of unset variables; $[m] = pn$, and $\mathbb{P}[|m - [m]| > n^{2/3}] \le e^{-\Omega(n^{1/3})}$. Condition henceforth on $m \in [pn \pm n^{2/3}]$.

Each clause survives independently with prob. $q = (1-p)^3 \pm o(1)$. Let $M = \Theta(n)$ be the original number of clauses; then the residual clause count M' satisfies $M' = (q \pm o(1))M$ w.h.p.

Decision-tree influence bound. Any depth-T decision tree adaptively queries at most T variables. Reveal the m variables in a fixed order; define the Doob martingale for the satisfiability indicator $X \in \{0,1\}$. Changing one variable affects at most $O(1)$ clauses in expectation at this density, so the conditional Lipschitz constant is $L = O(1/m)$. Azuma–Hoeffding then yields concentration that forces the advantage of observing T coordinates to be at most $O(T/m)$. Setting $T = O((\log n)^d)$ and $m = \Theta(pn)$ gives advantage $o(1)$; take n large so that $O(T/m) \le 1/6$.

SAT-balance event $\mathcal{B}$. Let $\mathcal{B}$ be the event that $\left|\mathbb{P}[\mathbf{SAT}(\varphi\restriction\rho) = 1] - 1/2\right| \le 1/6$.

Parameter choice for Lemma 3.7. We fix $\epsilon = c_4 n^{-\alpha}$ with $\alpha > 0$ small enough to keep the rejection probability of the balance test below n^{-5}. The Azuma–Hoeffding bound is then applied conditionally on the balance event, yielding the claimed $1/2 - c_4 n^{-\Omega(1)}$ correlation gap.

Proof details for Lemma 3.7. We bound the rejection probability of the balance filter and apply the influence bound from Lemma 3.5 under the balance condition, as detailed above, to derive the stated correlation gap.

References

- J. Håstad. *Computational Limitations of Small-Depth Circuits.* MIT Press, 1987.
- M. Chalupa. Volume I - Bounds - Formal Limits of Computability. Zenodo, 2025. DOI: 10.5281/zenodo.16408248.
- M. Chalupa. Auditability Beyond Computation: A Formal Model of Structural Drift and Semantic Stability. Zenodo, 2025. DOI: 10.5281/zenodo.16600703.
- M. Chalupa. Proof Integrity: Structural Drift and Semantic Stability in Computational Complexity. Zenodo, 2025. DOI: 10.5281/zenodo.15872999.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il