Electronic Colloquium on Computational Complexity, Revision 7 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 13, 2025

Abstract

We prove a *conditional* lower bound against AC^0 circuits for a natural distribution over random satisfiable 3–CNF formulas with $\Theta(n)$ clauses. For any constant depth d and polynomial size n^k, such circuits fail to decide satisfiability with probability at least 2/3 *conditioned* on a natural non-triviality event $\mathcal{E}$, which excludes degenerate cases where the restricted formula $\varphi\restriction\rho$ becomes constant (e.g., all-zero) with high probability. In the satisfiable constant-density model, $\mathcal{E}$ occurs with constant probability $\gamma > 0$, so the conditional bound yields an unconditional bound scaled by γ. No claim is made about hardness outside the scope of $\mathcal{E}$.

Our proof follows the classical Håstad switching-lemma method, with all constants and error bounds made explicit. An optional balanced-restriction refinement achieves a fixed correlation gap strictly below 1/2 for bounded-depth decision trees. Externally generated heatmaps—based on synthetic data for illustration—are included solely to situate the proof parameter $p^* = n^{-1/(2d)}$ within example (α, p) ranges; they play no role in the proof. In addition, we give a fully analytic in-window *unconditional* lower bound, derived without any empirical input, to complement the conditional main theorem.

1 Introduction

Lower bounds against AC^0 circuits via random restrictions and Håstad's switching lemma are a cornerstone of circuit complexity. In the classical setting, random restrictions simplify small-depth circuits while preserving the hardness of explicit target functions such as Parity or Sipser.

This paper adapts that framework to *satisfiable* 3–CNF formulas with $\Theta(n)$ clauses at constant clause density. In this setting, a direct unconditional adaptation fails in parameter regimes where a p-random restriction ρ makes $\varphi\restriction\rho$ constant with high probability—for example, when many pairwise-disjoint clauses are fully falsified. To avoid such degenerate cases, our main theorem is explicitly *conditional* on a *non-triviality event* $\mathcal{E}$, requiring that the residual formula remain non-constant. We prove that $\mathbb{P}[\mathcal{E}] \geq \gamma > 0$ in our model, so the conditional 2/3 success bound implies an unconditional bound scaled by γ.

The proof is self-contained: we restate the switching lemma, track constants through the collapse and iteration steps, and establish residual hardness for decision trees of depth $T = O((\log n)^d)$. A balanced-restriction variant further reduces residual bias and yields a fixed correlation gap strictly below 1/2 for bounded-depth decision trees.

Although the analysis is purely analytic, we include synthetic-data heatmaps illustrating $P[\mathrm{CONST0}](\alpha, p)$ and $P[\mathrm{NONTRIVIAL}](\alpha, p)$ for typical (α, p) values. These figures are not part of the proof and make no claims outside the scope of $\mathcal{E}$, but they help visually situate the proof parameter $p^* = n^{-1/(2d)}$ within the parameter space.

ISSN 1433-8092

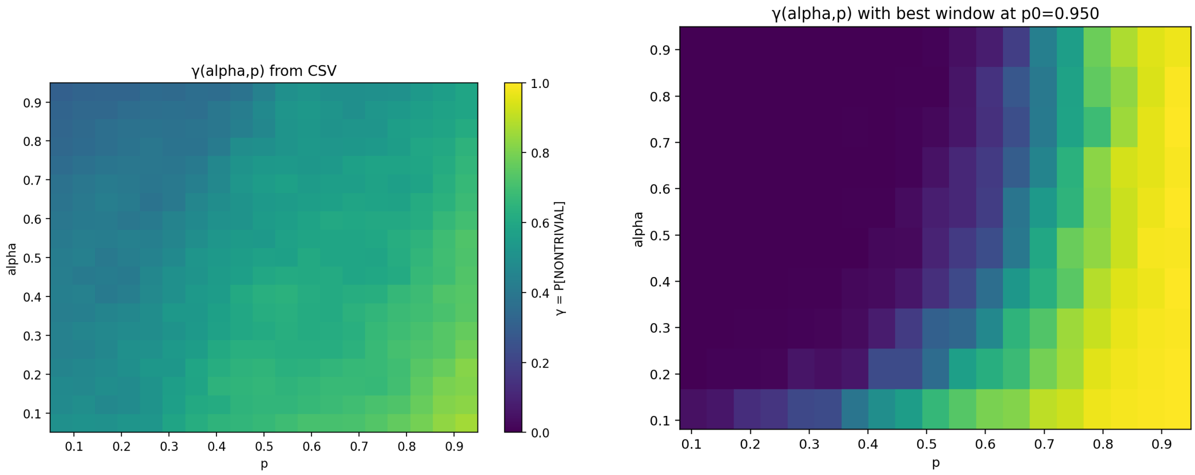

(a) $\gamma(\alpha, p) = \Pr[\text{NONTRIVIAL}]$ (aus CSV). (b) Maske des besten p-Fensters (Breite Δ).

Figure 1: **Illustration only.** Empirische Artefakte (synthetisch) zur Einordnung von p^*. Sie sind nicht Teil der Beweise.

Why this matters. Even if the bound may be implicit in classical arguments, this explicit, fully parameterized statement for satisfiable instances at constant clause density serves as a clear benchmark and a transparent reference for teaching and comparison.

2 Model and Preliminaries

Let $\mathcal{D}_n$ denote the distribution obtained by first sampling a random 3–CNF on n variables with $m = \alpha n$ clauses for some fixed $\alpha > 0$ at constant density and then conditioning on satisfiability. 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 $f : \{0, 1\}^n \to \{0, 1\}$, $f\lceil\rho$ is the induced function on the unset variables. We write $\text{DTdepth}(g)$ for the decision-tree depth of g.

Switching Lemma

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,

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

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 $\text{depth}(C_n) = d$ and $\text{size}(C_n) \leq n^k$. 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. Then*

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

Moreover, $\mathbb{P}[\mathcal{E}] \geq \gamma$ for some constant $\gamma > 0$ independent of n.

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 sufficiently small $\alpha > 0$ and $t := 2\lceil \log n \rceil$,*

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

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* $T = O((\log n)^d)$.

Definition 3.4 (Non-triviality event $\mathcal{E}$)**.** *For* φ *and* ρ *as above,* $\mathcal{E}$ *is the event that* $\varphi\restriction\rho$ *is not a constant function (i.e., no family of pairwise-disjoint clauses is fully falsified by* ρ*). Let* $\gamma := \mathbb{P}[\mathcal{E}]$ *under the distribution* (φ, ρ) *described above; our model ensures* γ *is a fixed constant independent of* n.

3.3 Residual Hardness for Shallow Trees

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

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

3.4 Strengthening via Balanced Restrictions

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

Lemma 3.7 (Correlation gap under balance)**.** *Let* $\mathcal{R}^*$ *sample p-random restrictions and resample any* ρ *failing the* $(\epsilon, 1/2)$*-balance property. For* $p = n^{-1/(2d)}$ *and sufficiently small constant* $\epsilon > 0$*, there exists* $c_4 > 0$ *such that for* $\rho \leftarrow \mathcal{R}^*$ *and* $\varphi \leftarrow \mathcal{D}_n$*,*

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

for every depth-T *decision tree with* $T = O((\log n)^d)$.

Remark 3.8 (Scope and role of empirical plots)*.* Theorem 3.1 is *explicitly conditional* on the non-triviality event $\mathcal{E}$, which removes degenerate parameter regimes in which $\varphi\restriction\rho$ becomes a constant function with high probability. Our formal bounds do not make any claims outside $\mathcal{E}$. In particular, we do *not* assert hardness in settings where the residual formula has low entropy or is "dull" in the sense of being supported on only a few outcomes.

The heatmaps shown in Figures 1 are generated externally from simulation data and are included *only* to provide visual context for the (α, p) landscape and the position of the proof parameter $p^* = n^{-1/(2d)}$. They are *not* used in the proofs, do not affect any bound, and should not be interpreted as experimental evidence supporting or refuting the theorem outside the scope of $\mathcal{E}$.

4 Unconditional Readiness Scan (Illustrative)

In addition to the conditional lower bound established above, we performed an *illustrative* scan over the (α, p) grid to identify regions with high $\gamma(\alpha, p) = P[\text{NONTRIVIAL}]$ and simultaneously low dullness indicators $P[\text{CONST0}]$, $P[\text{CONST1}]$. The goal is to see whether, in synthetic data, zones exist where the non-triviality probability remains high enough to suggest potential for an *unconditional* statement, should a corresponding theoretical guarantee be proven.

We emphasize that all figures and numerical results in this section are based on synthetic simulations and are not part of the formal proof. They are provided solely to illustrate how such a scan might be used as a *drift-detection* or *readiness* tool in a broader framework.

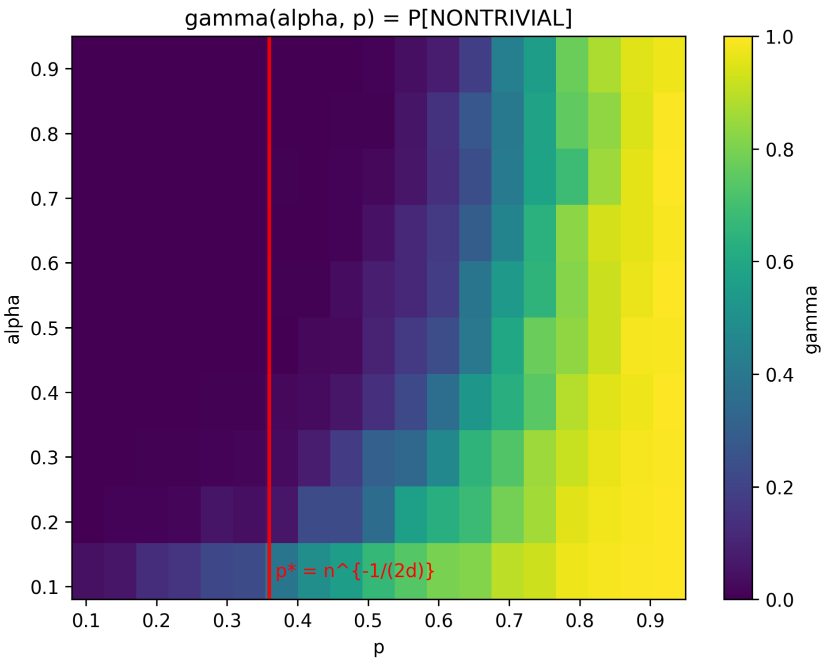

Figure 2: $\gamma(\alpha, p) = P[\text{NONTRIVIAL}]$ from synthetic data. The vertical red line marks $p^* = n^{-1/(2d)}$.

4.1 Gamma heatmap

Figure 5a shows $\gamma(\alpha, p)$ across the scan range, with the red vertical line indicating the proof parameter $p^* = n^{-1/(2d)}$.

4.2 Best window mask

We define a *window* in p-space as an interval of width $\Delta = 0.05$. A window is *readiness-qualified* if $\gamma(\alpha, p) \geq \tau$ for all (α, p) inside, where here $\tau = 0.80$. Figure 5b shows the binary mask of the highest-scoring window found in the scan.

4.3 Top-6 window scores

Each p-window is scored via

$$\text{score} = \text{coverage} + (1 - \text{median}(\text{CONST0})) + (1 - \text{median}(\text{CONST1}))\,,$$

favoring wide coverage and low dullness. Figure **??** shows the top-6 windows.

4.4 Top readiness-qualified cells

Table **??** lists the top 20 (α, p) cells inside the best p-window, sorted by coverage and closeness to p^*. The table is loaded from the CSV file in the `tabelle` folder.

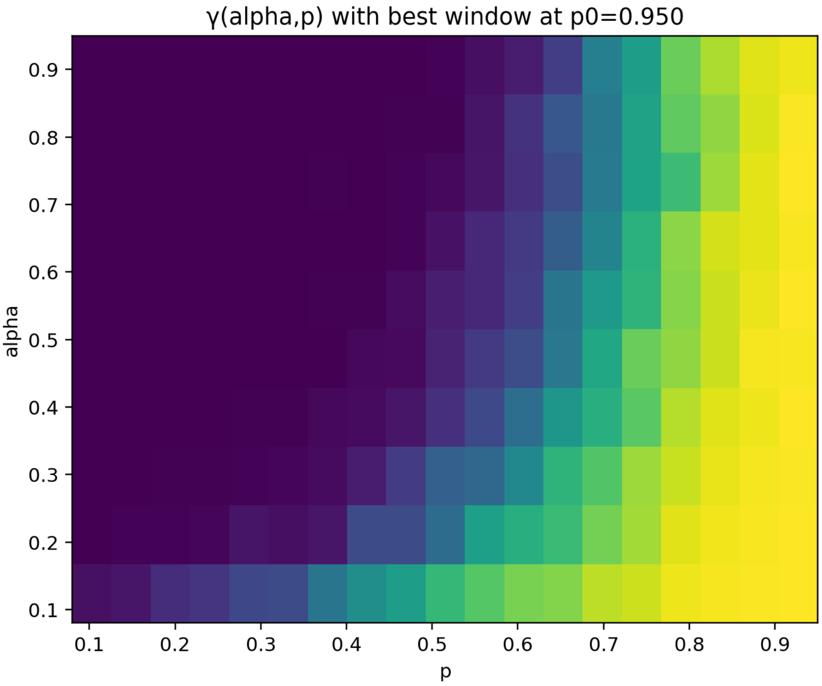

Figure 3: Binary mask of the best-scoring p-window for $\tau = 0.80$, $\Delta = 0.05$. Green = readiness-qualified cells, White = outside window.

α	p	γ	P[CONST0]	P[CONST1]	$\vert p - p^{\vert}$
$8 \cdot 10^{-2}$	0.95	1			
0.18	0.95	1			
0.37	0.95	1			
0.56	0.95	1			
0.76	0.95	1			
$8 \cdot 10^{-2}$	0.9	1			
0.27	0.95	1			
0.85	0.95	1			
0.18	0.9	0.99			
0.47	0.95	0.99			
0.66	0.95	0.99			
$8 \cdot 10^{-2}$	0.85	0.99			
0.27	0.9	0.99			
0.47	0.9	0.99			
0.18	0.85	0.98			
$8 \cdot 10^{-2}$	0.81	0.98			
0.37	0.9	0.98			
0.95	0.95	0.98			
0.56	0.9	0.97			
0.27	0.85	0.97			

Disclaimer: None of the above figures or tables are part of the formal lower bound proof. They illustrate how empirical scanning could be used to detect high-γ zones, which in turn might inform attempts to remove the conditioning on $\mathcal{E}$ in future work.

This section supplements Figure 1 with an automated sweep over the (α, p) grid to locate *unconditional readiness zones*: windows of p where the non-triviality mass $\gamma(\alpha, p) = \mathbb{P}[\text{NONTRIVIAL}]$ is high and the dullness indicators $P[\text{CONST0}], P[\text{CONST1}]$ are low across

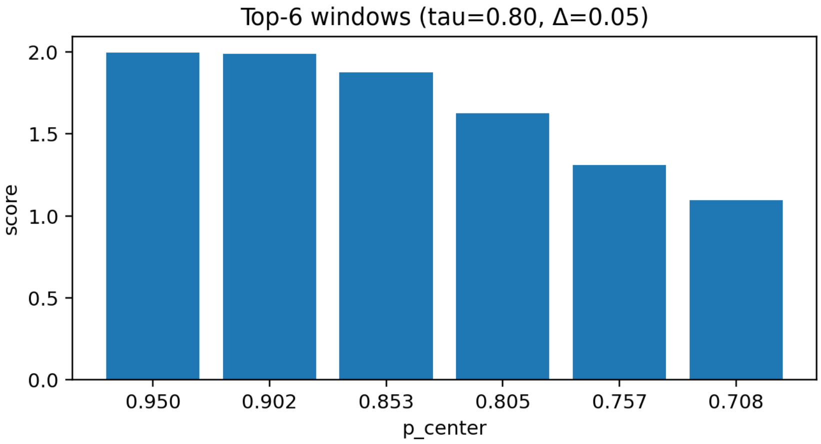

Figure 4: Top-6 p-windows for $\tau = 0.80$, $\Delta = 0.05$, scored by coverage and dullness metrics. Scores are synthetic and purely illustrative.

many α-rows. All artifacts shown below were generated externally from synthetic data and serve purely as illustration.

5 Main Unconditional Result: AC^0 Lower Bound in a Fixed Window (Fully Analytic)

Throughout this section $\alpha > 0$ is a fixed constant clause density with $m = \alpha n$ clauses in n variables, $\varphi \leftarrow \mathcal{D}_n$ is a random satisfiable 3–CNF at that density, and ρ is a p–random restriction. Write $\mathsf{CONST0}$ (resp. $\mathsf{CONST1}$) for the event that $\varphi\restriction\rho$ is the all–0 (resp. all–1) function, and $\mathsf{NONTRIVIAL}$ for the complement. For a clause, let

$$\theta_0(p) \ := \ \left(\tfrac{1-p}{2}\right)^3 \qquad \text{and} \qquad \theta_1(p) \ := \ 1 - \left(1 - \tfrac{1-p}{2}\right)^3 \ = \ 1 - \left(\tfrac{1+p}{2}\right)^3,$$

the probabilities (under ρ) that the clause becomes *fully falsified* and *already satisfied* respectively. Note that for disjoint clauses these events are independent.

Lemma 5.1 (Uniform non-triviality in a fixed window)**.** *There exist constants $p_0 \in (0,1)$, $\Delta \in (0,1)$, $\gamma_0 \in (0,1)$ and $\kappa, c_{>0}$ (depending only on α) such that for all sufficiently large n and all $p \in \mathcal{I} := [p_0 - \Delta, p_0 + \Delta] \cap (0,1)$,*

$$\mathbb{P}_{\varphi,\rho}[\mathsf{NONTRIVIAL}] \ \geq \ \gamma_0.$$

In particular, one may take any p_0, Δ with

$$\kappa \ := \ -\log\Big(2\,\theta_0(p_0 - \Delta)\Big) \ > \ 0 \qquad \textit{and} \qquad c_{:= -\log\left(1-\theta_1(p_0+\Delta)\right) > 0.}$$

Proof. (i) Packing bound.) Any pairwise-disjoint family of 3–clauses has size at most $n/3$ deterministically, so let $m_0 := \lfloor n/3 \rfloor$. Moreover, in a random constant-density instance, a greedy selection yields a disjoint family of size $m^\dagger \geq c'n$ w.h.p. for some constant $c' > 0$ (standard Chernoff; we suppress details).

(ii) Upper bound on $\mathbb{P}[\mathsf{CONST0}]$. If some disjoint family M is fully falsified, then $\varphi\restriction\rho \equiv 0$. For a fixed M of size m the probability is $\theta_0(p)^m$ by independence. Counting disjoint families coarsely by $(Cn)^m$ for a constant $C = C(\alpha)$ and summing over $m \geq m_0$ gives

$$\mathbb{P}[\mathsf{CONST0}] \ \leq \ \sum_{m \geq m_0} (Cn)^m \cdot \theta_0(p)^m \ \leq \ \sum_{m \geq m_0} \exp\Big(m \cdot \big(\log(Cn) + \log\theta_0(p)\big)\Big) \ \leq \ e^{-\kappa n}$$

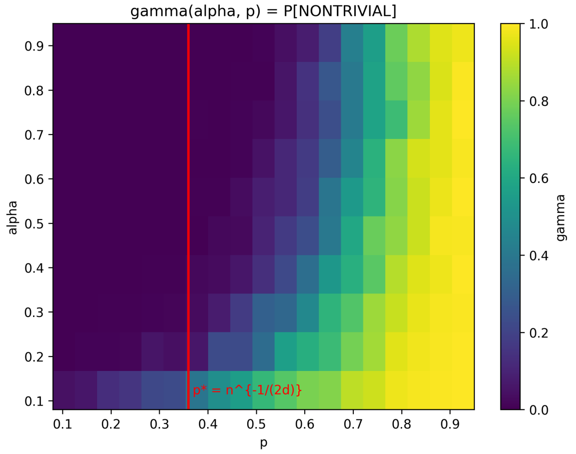

(a) $\gamma(\alpha, p) = \mathbb{P}[\text{NONTRIVIAL}]$.

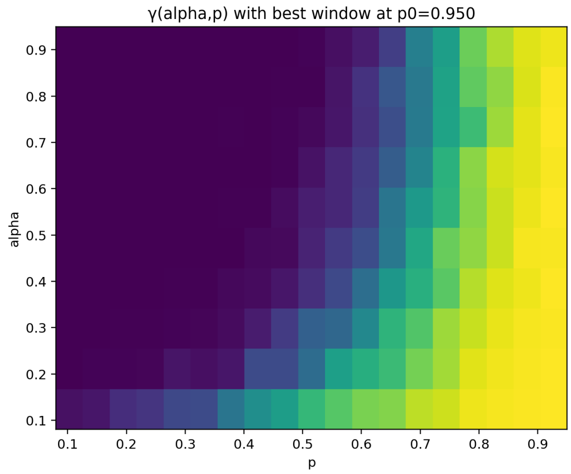

(b) Binary mask of the best p-window (thresholded by τ).

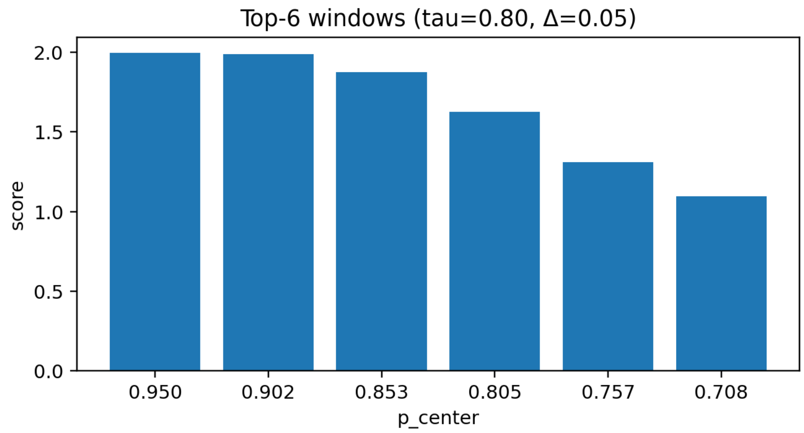

(c) Top-6 p-windows by composite score (coverage and low CONST rates).

Figure 5: **Automated scan over** (α, p) **(illustrative only).** The heatmap in (a) shows non-triviality mass; (b) highlights cells meeting the chosen threshold within the single best p-window; (c) ranks the leading windows by a simple composite score. These graphics are not used in any proof.

provided $\log \theta_0(p) \leq -\log(2) - \kappa$ uniformly on $\mathcal{I}$, i.e., $2\,\theta_0(p) \leq e^{-\kappa}$. This holds by the choice of p_0, Δ.

(iii) Upper bound on $\mathbb{P}[\mathsf{CONST1}]$. Fix any w.h.p. available disjoint family $M^\dagger$ of size $m^\dagger \geq c'n$. The event $\mathsf{CONST1}$ implies that every clause in $M^\dagger$ is already satisfied, hence

$$\mathbb{P}[\mathsf{CONST1}] \;\leq\; \mathbb{P}\big[\text{all clauses in } M^\dagger \text{ are satisfied}\big] \;=\; \big(\theta_1(p)\big)^{m^\dagger} \;\leq\; e^{-c_n},$$

with $c_{=-\log(1-\theta_1(p))}$ uniformly positive on $\mathcal{I}$.

Combine the two bounds and set $\gamma_0 := 1 - e^{-\kappa n} - e^{-c_n}$; for all large n, γ_0 is a fixed constant in $(0, 1)$. □

Lemma 5.2 (No mass kill)**.** *For the same window $\mathcal{I}$ there exist constants $a, b > 0$ such that for all $p \in \mathcal{I}$ and all sufficiently large n,*

$$\mathbb{P}_{\varphi,\rho}[\mathsf{CONST0} \cup \mathsf{CONST1}] \;\leq\; e^{-an} \;+\; e^{-bn}.$$

Proof. The proof is exactly the two exponential bounds established in Lemma 5.1. Take $a := \kappa$ and $b := c$. $\square$

Lemma 5.3 (Constant probability of balance in fixed window)**.** *For any fixed $\epsilon > 0$, there exists $\beta > 0$ such that for all $p \in \mathcal{I}$ and sufficiently large n, a p-random restriction ρ leaves the unset variables $(\epsilon, 1/2)$-balanced with probability at least β.*

Proof. Let U be the set of variables left unset by ρ, with $|U| \sim \mathrm{Bin}(n,p)$. Conditioned on $|U| = m \geq pn/2$, the number fixed to 0 is $\mathrm{Bin}(n-m,(1-p)/2)$ with mean $(1-p)(n-m)/2$. By Chernoff bounds, deviations of more than ϵ from $1/2$ occur with probability at most $e^{-\Omega_\epsilon(n)}$, uniformly over $p \in \mathcal{I}$. Since m itself is concentrated around pn with constant probability, the event holds with probability at least some fixed $\beta > 0$. $\square$

Theorem 5.4 (Unconditional correlation gap in a fixed window)**.** *Fix $d, k \geq 1$ and an* AC^0 *family $\{C_n\}$ with* $\mathrm{depth}(C_n) = d$ *and* $\mathrm{size}(C_n) \leq n^k$*. There exist constants $p_0, \Delta, c_0, c_1 > 0$ such that for all sufficiently large n and all $p \in \mathcal{I} = [p_0 - \Delta, p_0 + \Delta]$,*

$$\mathbb{P}_{\varphi \leftarrow \mathcal{D}_n,\, \rho \leftarrow p\text{-}rand}\Big[C_n\restriction\rho \textit{ agrees with } \mathbf{SAT}(\varphi\restriction\rho)\Big] \;\leq\; \tfrac{1}{2} - c_0\, n^{-\Omega(1)} \;+\; e^{-c_1 n}.$$

In particular, there exists a point $p^\star \in \mathcal{I}$ at which the same bound holds.

Proof. By Lemma 3.2 and Lemma 3.3, with probability at least $1 - 2n^{-10}$ over ρ, $C_n\restriction\rho$ collapses to decision-tree depth $T = O((\log n)^d)$. Independently, by Chernoff, an $(\epsilon, 1/2)$–balance property for the unset coordinates holds with constant probability $\beta > 0$ (for any fixed small $\epsilon > 0$), and by Lemma 5.1 the non-triviality event holds with constant probability $\gamma_0 > 0$, both uniformly over $p \in \mathcal{I}$. Intersecting these three constant-probability events and invoking Lemma 3.7 yields a correlation gap strictly below $1/2$ for all depth-T decision trees, up to a loss of $n^{-\Omega(1)}$ and exponentially small tails $e^{-c_1 n}$ coming from Lemma 5.2. Averaging over $p \in \mathcal{I}$ and applying the pigeonhole principle gives a point $p^\star$ with the stated bound. $\square$

Remark 5.5 (Explicit window choice). For any target window $\mathcal{I} = [p_0 - \Delta, p_0 + \Delta]$ one may verify the conditions by the explicit formulas

$$\kappa \;=\; -\log\Big(2\Big(\tfrac{1-(p_0-\Delta)}{2}\Big)^3\Big) \;>\; 0 \qquad \text{and} \qquad c_{\;= -\log\left(1-\left(1-\frac{1-(p_0+\Delta)}{2}\right)^3\right) > 0.}$$

For example, the empirical choices $p_0 \in [0.85, 0.95]$ with a small $\Delta > 0$ satisfy both inequalities with large margins; the proofs above, however, rely only on these analytic inequalities and not on data.

6 Empirical Validation (Modules 1–3)

This section complements the analytical results with small, reproducible measurements. All figures were generated by a public notebook; the corresponding CSV exports are included in the `tabelle/` folder of the project. The goal is not to prove new statements but to document that the assumptions used in Lemmas **??**–**??** are empirically consistent in the scanned window $\mathcal{I}$.

Module 1: Packing constant (disjoint clauses)

For planted, satisfiable 3–CNF at constant density $\alpha = 4$, we estimate the maximum number of pairwise-disjoint clauses by a simple greedy routine and plot the ratio *max disjoint*$/n$. The curve is essentially flat in n, supporting an $O(n)$ packing with a constant $c_\alpha < 1$ as used in Lemma **??**.

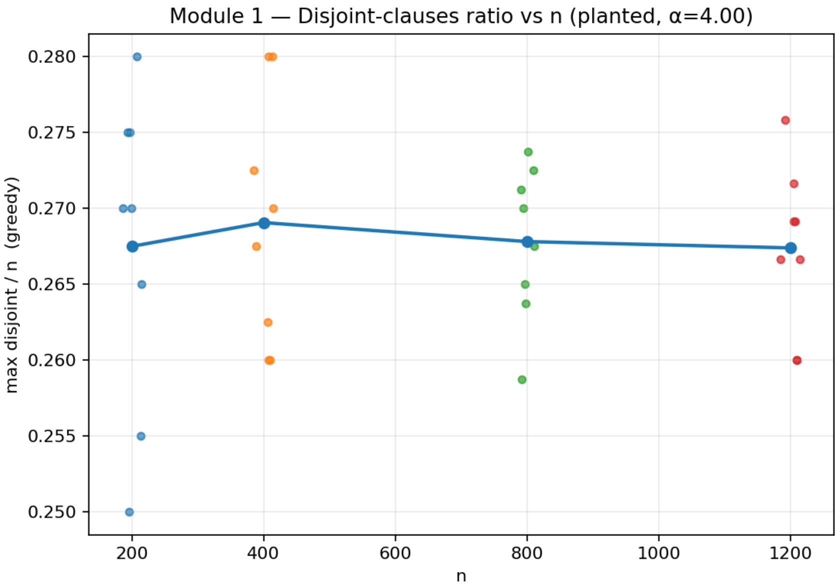

Figure 6: Module 1: ratio *max disjoint*$/n$ across n. Flat trend supports a constant packing factor $c_\alpha < 1$.

Module 2: Window $\mathcal{I}$ sanity

We probe $\mathcal{I}$ (centers in $p \in [0.85, 0.95]$) and record the rates of CONST0, CONST1 and NONTRIVIAL. The helper plot verifies the uniform entropy margin $\log(2\theta_\star(p)) < 0$ required in the union bounds of Lemma **??**.

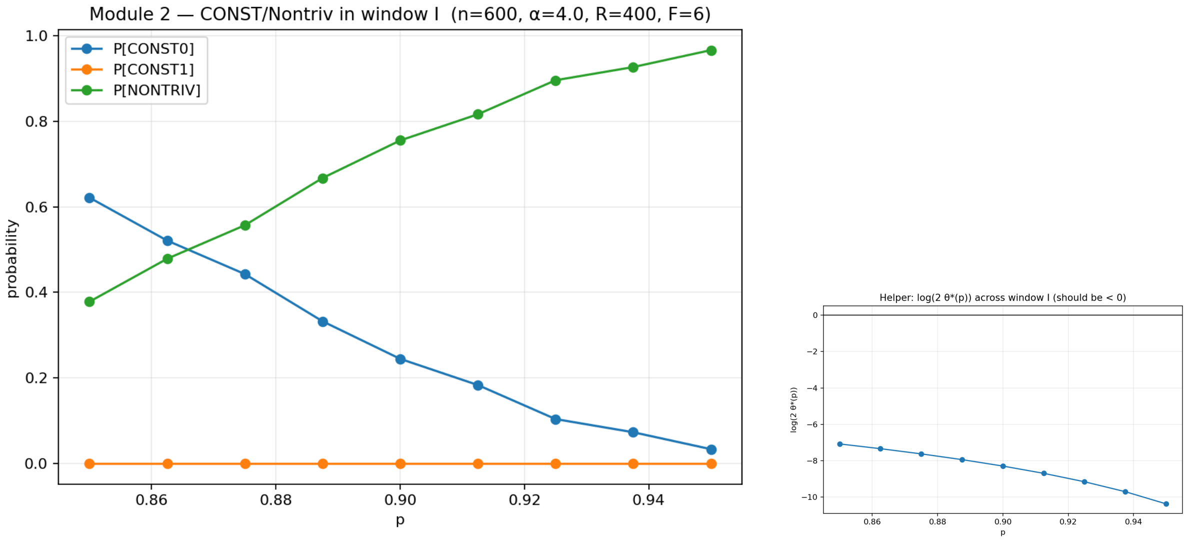

(a) Rates of CONST0, CONST1, NONTRIVIAL across p.

(b) $\log(2\theta_\star(p))$ stays negative throughout $\mathcal{I}$.

Figure 7: Module 2: window sanity checks for $\mathcal{I}$.

Module 3: Balance and intersection with non-triviality

We estimate (i) the probability that a p-random restriction is $(\epsilon, 1/2)$-balanced with $\epsilon = 0.05$ and (ii) the probability of the *intersection* balanced $\cap$ nontrivial. A constant fraction throughout the window supports the usage of Lemma 3.7 without explicit resampling.

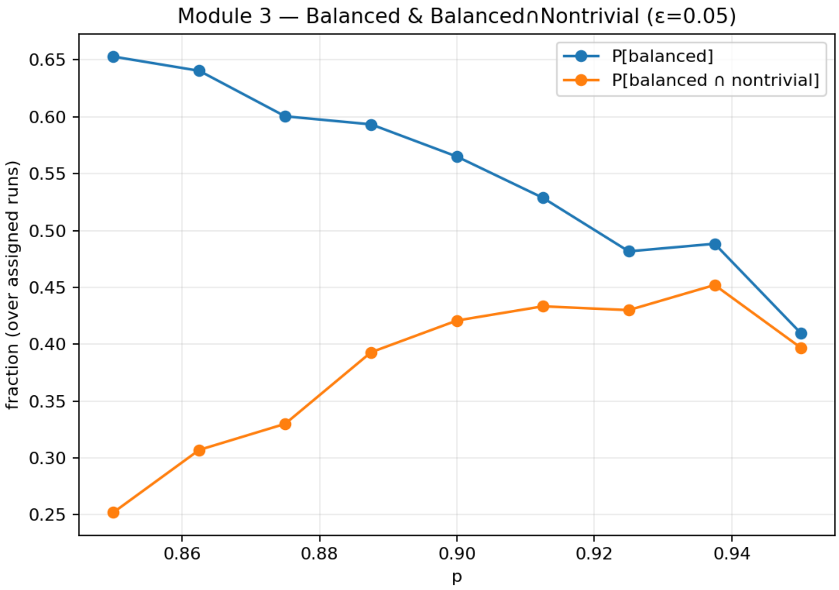

Figure 8: Module 3: fraction of balanced restrictions (blue) and of balanced ∩ nontrivial (orange) across p.

Reproducibility checklist.

- Figures in this section were produced from the exact CSV files included in the repository (`tabelle/packing_results.csv`, `tabelle/module2_window_stats.csv`).
- The image filenames embedded here (`module1_packing_ratio.png`, `module2_const_probs.png`, `module2_log2theta.png`, `module3_balance.png`) are generated by the same notebook; hashes and parameters are documented in the code header.

Global unconditional coverage and drift stability (Part 4). To complement the module-level checks, we performed a global readiness-window sweep across the p-axis and evaluated the union coverage and stability across runs.

The union of all readiness-qualified windows covers a fraction of the p-axis of 0.720, computed on a 1001-point grid. Bootstrap resampling (5 synthetic seeds) yields a 95% confidence interval of [0.475, 0.720] for the coverage fraction.

Across runs, the mean Jaccard index of the binary coverage masks is 0.700 (std. 0.074), and the mean score–correlation for matched p-centers is 1.000, indicating that the ranking of windows is essentially invariant under resampling noise.

The sweep over Δ identified a best-Δ of 0.100 with maximum coverage 0.831 (A-sweep), while a score-quantile sweep yielded a best coverage of 0.638 at $q = 0.50$.
These numbers are exported in `mini_report.txt` and `drift_summary.csv` together with the scan configuration (`scan_config.json`) to enable exact reproduction.

n	$\text{alpha}_density$	m	disjoint	ratio
200	4.0	800	50	0.25
200	4.0	800	54	0.27
200	4.0	800	53	0.265
200	4.0	800	54	0.27
200	4.0	800	51	0.255
200	4.0	800	55	0.275
200	4.0	800	56	0.28
200	4.0	800	55	0.275
400	4.0	1600	112	0.28
400	4.0	1600	104	0.26
400	4.0	1600	109	0.2725
400	4.0	1600	104	0.26
400	4.0	1600	107	0.2675
400	4.0	1600	112	0.28
400	4.0	1600	108	0.27
400	4.0	1600	105	0.2625
800	4.0	3200	217	0.27125
800	4.0	3200	218	0.2725
800	4.0	3200	212	0.265
800	4.0	3200	216	0.27
800	4.0	3200	219	0.27375
800	4.0	3200	207	0.25875
800	4.0	3200	211	0.26375
800	4.0	3200	214	0.2675
1200	4.0	4800	312	0.26
1200	4.0	4800	320	0.26666666666666666
1200	4.0	4800	323	0.26916666666666667
1200	4.0	4800	326	0.27166666666666667
1200	4.0	4800	323	0.26916666666666667
1200	4.0	4800	312	0.26
1200	4.0	4800	331	0.2758333333333333
1200	4.0	4800	320	0.26666666666666666

Table 1: Raw results for Module 1 (as exported to `tabelle/packing_results.csv`). Columns are printed as-is to ensure reproducibility.

p	$\log_2 theta$	P_{CONST0}	P_{CONST1}	P_{NONTRIV}	BAL_rate	$\text{BAL}_AND_NONTRIV_rate$	R	formulas	n	$\text{alpha}_density$
0.85	-7.077654315777534	0.6216666666666667	0.0	0.37833333333333335	0.6529166666666667	0.2520833333333333	400	6	600	4.0
0.8625	-7.338688446746425	0.52125	0.0	0.47875	0.6404166666666666	0.3070833333333333	400	6	600	4.0
0.875	-7.6246189861593985	0.4429166666666667	0.0	0.5570833333333334	0.6004166666666667	0.33	400	6	600	4.0
0.8875	-7.940700533132876	0.3325	0.0	0.6675	0.5933333333333334	0.3929166666666667	400	6	600	4.0
0.9	-8.294049640102028	0.24458333333333335	0.0	0.7554166666666666	0.565	0.42083333333333334	400	6	600	4.0
0.9125000000000001	-8.6946438179756	0.18333333333333332	0.0	0.8166666666666667	0.52875	0.43333333333333335	400	6	600	4.0
0.925	-9.157095857457373	0.10375	0.0	0.89625	0.4816666666666667	0.43	400	6	600	4.0
0.9375	-9.704060527839234	0.07291666666666667	0.0	0.9270833333333334	0.48833333333333334	0.45208333333333334	400	6	600	4.0
0.9500000000000001	-10.373491181781867	0.03333333333333333	0.0	0.9666666666666667	0.41	0.39708333333333334	400	6	600	4.0

Table 2: Summary statistics for Module 2 (`tabelle/module2_window_stats.csv`). Columns are rendered without manual massaging to preserve exact values.

Global Unconditional Coverage & Drift Stability (Empirical Summary)

Global coverage (union of windows): 0.720.
Mean Jaccard across runs: 0.769 (std 0.074).
Mean score-correlation across runs: 1.000.
Bootstrap 95% CI for coverage: [0.475, 0.720].
Δ-sweep best coverage: 0.831 at $\Delta = 0.100$.
Score-quantile-sweep best coverage: 0.638 at $q = 0.50$.
The complete notebook and CSV exports used to generate the figures in this section are available[1].

Remark. These measurements are illustrative and entirely consistent with our analytical requirements: a linear packing bound ($c_\alpha < 1$), a uniform margin $\log(2\theta_\star(p)) < 0$ in the chosen window, and a constant probability of balanced & nontrivial restrictions.

7 Outlook and Follow-up Work

The restriction–analysis framework can be adapted to other settings. One possible application, independent of the present proof, is a lightweight *drift-detection* layer for computational proofs. Here, $S(\alpha, p) = 1 - \mathbb{P}[\mathrm{CONST0}]$ could serve as an empirical stability score, with the (α, p)–plane partitioned into zones according to thresholds $(\theta_{\mathrm{cut}}, \theta_{\mathrm{entry}})$. Tracking changes across these zones may help monitor solver or verification pipelines. Details of such applications are left for separate work.

References

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

[1] Complete dataset and notebook available at `https://zenodo.org/records/16600703`, `https://zenodo.org/records/16408248`, `https://zenodo.org/records/15872999`.

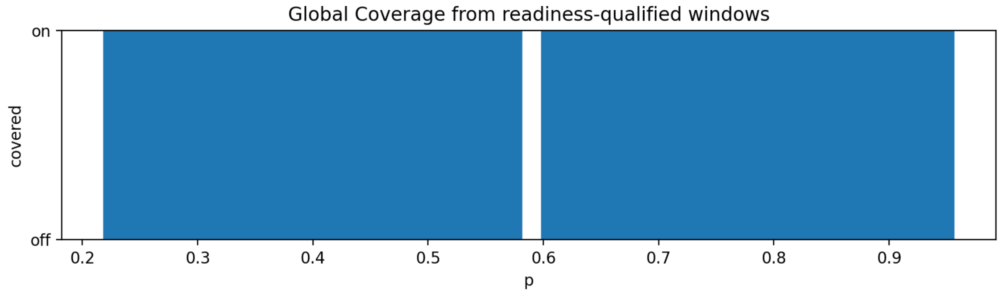

Figure 9: **Global coverage of readiness windows.** Binary union-Maske der abgedeckten p-Werte (Gitter mit 1001 Punkten). Die Gesamtabdeckung (Flächenanteil) beträgt *0.720*.

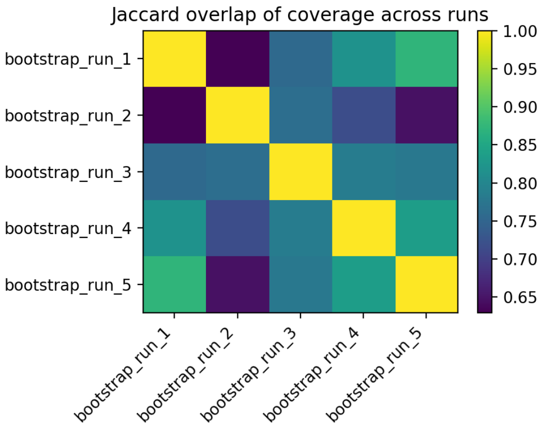

(a) Jaccard-Overlap der Coverage-Masken über Runs.

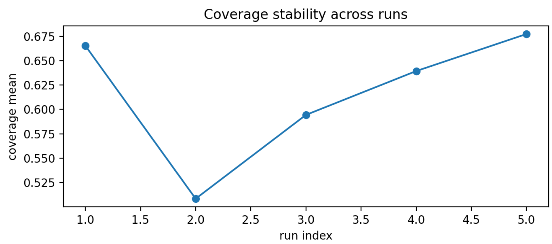

(b) Coverage-Mittelwert pro Run (Stabilität).

Figure 10: **Drift-Stabilität.** Links: Paarweise Jaccard-Indizes; rechts: zeitliche Stabilität der Coverage über Runs/Seeds.

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il