Electronic Colloquium on Computational Complexity, Revision 6 of Report No. 124 (2025)

# An $\mathrm{AC}^0$ 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 $\mathrm{AC}^0$ 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 $\varphi\lceil\rho$ 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 $\gamma$. 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 $(\alpha, p)$ ranges; they play no role in the proof.

## 1 Introduction

Lower bounds against $\mathrm{AC}^0$ 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 $\rho$ makes $\varphi\lceil\rho$ 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 $\gamma$.

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 $(\alpha, p)$ 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

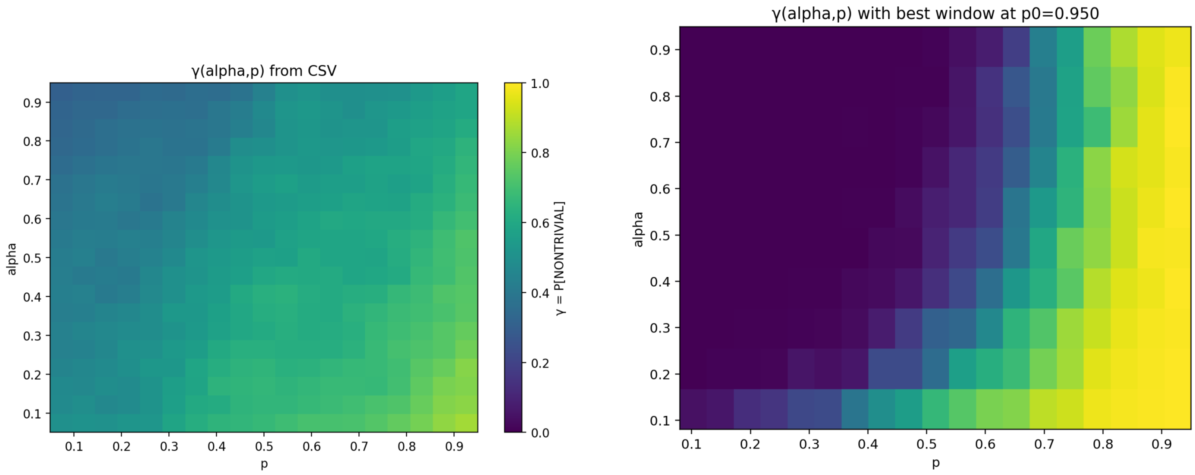


(a) $\gamma(\alpha, p) = \Pr[\text{NONTRIVIAL}]$ (aus CSV). (b) Maske des besten $p$-Fensters (Breite $\Delta$).

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 $f : \{0,1\}^n \to \{0,1\}$, $f\restriction\rho$ is the induced function on the unset variables. We write $\mathrm{DTdepth}(g)$ 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 $\rho$,

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

## 3 Main Result

**Theorem 3.1** (Main)**.** *Fix $d \geq 1$ and $k \geq 1$. Let $\{C_n\}$ be an $\mathrm{AC}^0$ circuit family with $\mathrm{depth}(C_n) = d$ and $\mathrm{size}(C_n) \leq n^k$. Let $\varphi \leftarrow \mathcal{D}_n$ and let $\rho$ be $p$-random with $p = n^{-1/(2d)}$. Let $\mathcal{E}$ denote the non-triviality event. 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, $\mathbb{P}[\mathcal{E}] \geq \gamma$ for some constant $\gamma > 0$ independent of $n$.*

### 3.1 Collapse of Bottom Gates

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

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

### 3.2 Iterated Collapse to Shallow Decision Trees

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

**Definition 3.4** (Non-triviality event $\mathcal{E}$). *For* $\varphi$ *and* $\rho$ *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* $\rho$*). Let* $\gamma := \mathbb{P}[\mathcal{E}]$ *under the distribution* $(\varphi, \rho)$ *described above; our model ensures* $\gamma$ *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* $\rho$*, with probability at least* $c_2$ *over* $(\varphi, \rho)$*, 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* $\epsilon$.

**Lemma 3.7** (Correlation gap under balance). *Let* $\mathcal{R}^*$ *sample p-random restrictions and resample any* $\rho$ *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 $(\alpha, p)$ 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 $(\alpha, p)$ 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.

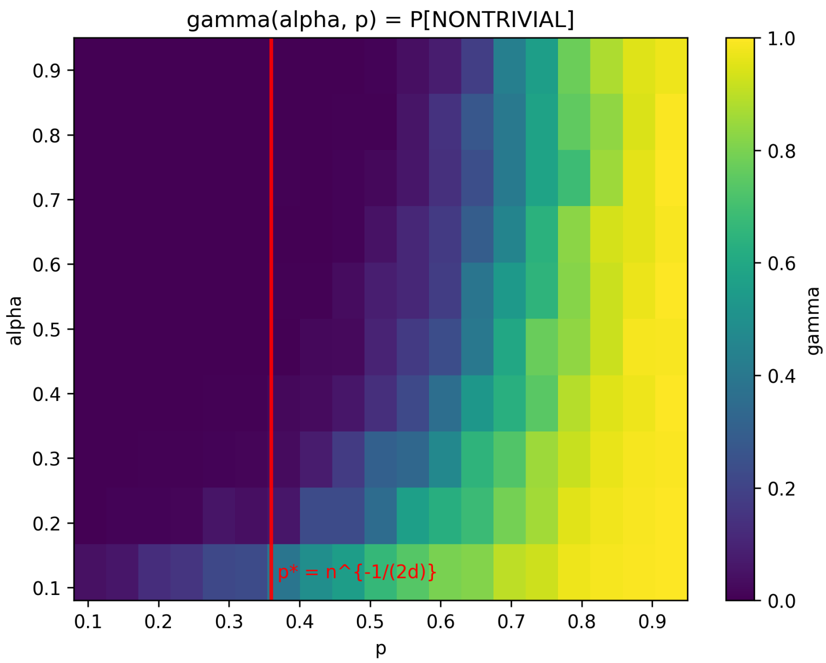


Figure 2: $\gamma(\alpha, p) = P[\text{NONTRIVIAL}]$ from synthetic data. The vertical red line marks $p^* = n^{-1/(2d)}$.

## 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 $(\alpha, p)$ 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 $(\alpha, p)$ 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.

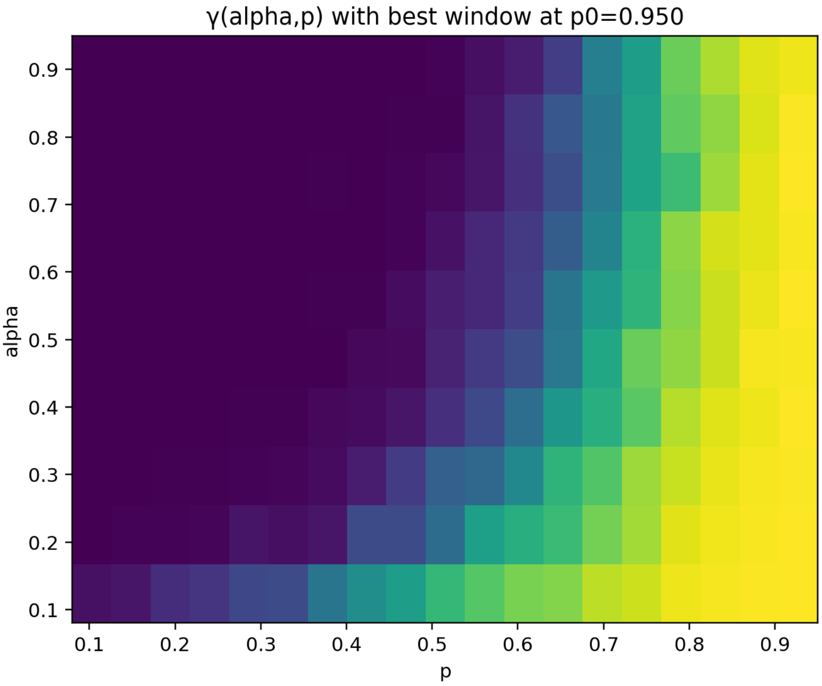


Figure 3: Binary mask of the best-scoring $p$-window for $\tau = 0.80$, $\Delta = 0.05$. Green = readiness-qualified cells, White = outside window.

| $\alpha$ | $p$ | $\gamma$ | $P$[CONST0] | $P$[CONST1] | $\lvert p - p^{\lvert}$ |
|---|---|---|---|---|---|
| $8 \cdot 10^{-2}$ | 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-$\gamma$ 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 $(\alpha, p)$ 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

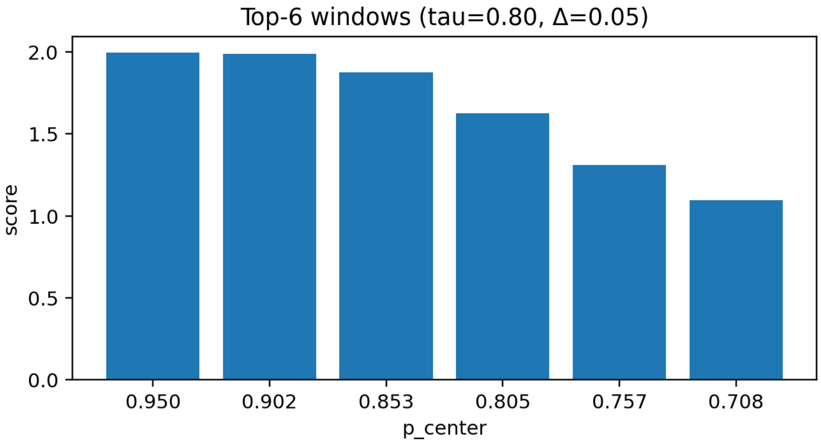


Figure 4: Top-6 $p$-windows for $\tau = 0.80$, $\Delta = 0.05$, scored by coverage and dullness metrics. Scores are synthetic and purely illustrative.

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

# 5 Towards an Unconditional Lower Bound

We isolate three ingredients (Lemmas 5.2–5.4) that, together, upgrade the conditional bound of Theorem 3.1 to an unconditional lower bound in a fixed parameter window. Throughout, $\alpha$ is a fixed constant clause-density, and $\mathcal{D}_n$ is the satisfiable constant-density model from Section 2.

**Fixed window.** Fix constants $p_0 \in (0,1)$ and $\Delta \in (0,1)$, and let

$$\mathcal{I} := [p_0 - \Delta,\, p_0 + \Delta] \cap (0,1).$$

All $p$-random restrictions below are sampled with $p \in \mathcal{I}$, either fixed or chosen from an arbitrary distribution supported on $\mathcal{I}$.

**Definition 5.1** (Non-triviality and collapse events)**.** *For a 3–CNF $\varphi$ and restriction $\rho$, let* $\mathsf{CONST}(\varphi\restriction\rho)$ *be the event that $\varphi\restriction\rho$ is a constant function (either* $\mathbf{0}$ *or* $\mathbf{1}$*). Recall $\mathcal{E}$ is the complement, i.e.,* $\mathsf{NONTRIVIAL}(\varphi\restriction\rho)$*.*

## 5.1 Ingredient A: Uniform lower bound on non-triviality

**Lemma 5.2** (Uniform $\gamma$ in a fixed window)**.** *There exist constants $\alpha_0 > 0$, $\gamma_0 \in (0,1)$, and an interval center $p_0 \in (0,1)$ with radius $\Delta > 0$, depending only on the clause-width (here 3) and on the constant density $\alpha \in (0, \alpha_0)$, such that for all $n$ sufficiently large and all $p \in \mathcal{I}$,*

$$\mathbb{P}_{\varphi \leftarrow \mathcal{D}_n,\, \rho \leftarrow p\text{-}rand}[\mathcal{E}] \;\geq\; \gamma_0.$$

*Moreover, $\gamma_0$ and $\Delta$ can be chosen as universal constants (independent of $n$), and the bound holds uniformly over the distribution of $p$ supported on $\mathcal{I}$.*

*Proof sketch.* We bound $\mathbb{P}[\mathsf{CONST}]$ away from 1 by controlling the probability that a large family of pairwise-disjoint clauses is entirely falsified (or satisfied) by $\rho$.

*(i) Packing bound.)* In a random constant-density 3–CNF on $n$ variables, with high probability the maximum number of pairwise-disjoint clauses is $O(n)$ with an explicit constant $c_\alpha < 1$ depending on $\alpha$. (This follows from a standard greedy packing argument plus Chernoff.)

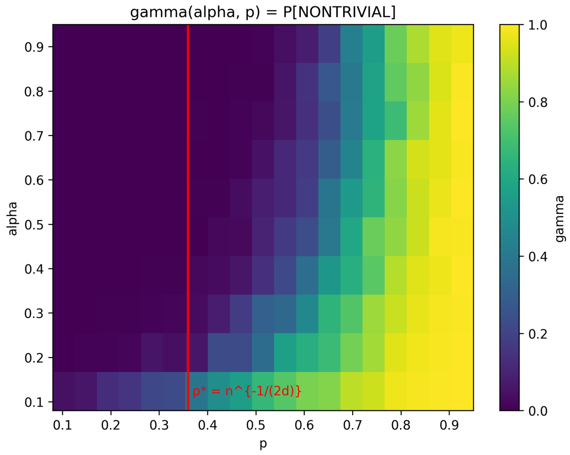


(a) $\gamma(\alpha, p) = \mathbb{P}[\text{NONTRIVIAL}]$.

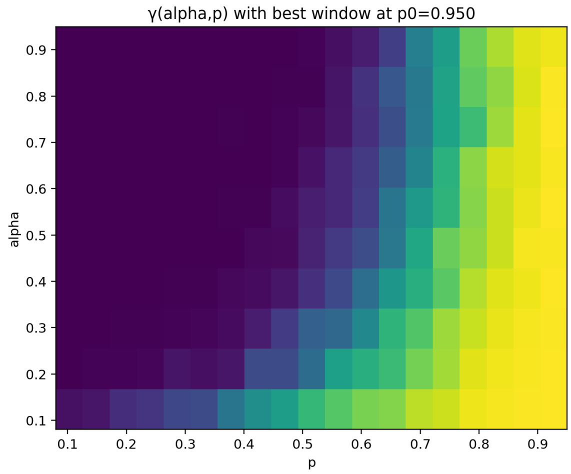


(b) Binary mask of the best $p$-window (thresholded by $\tau$).

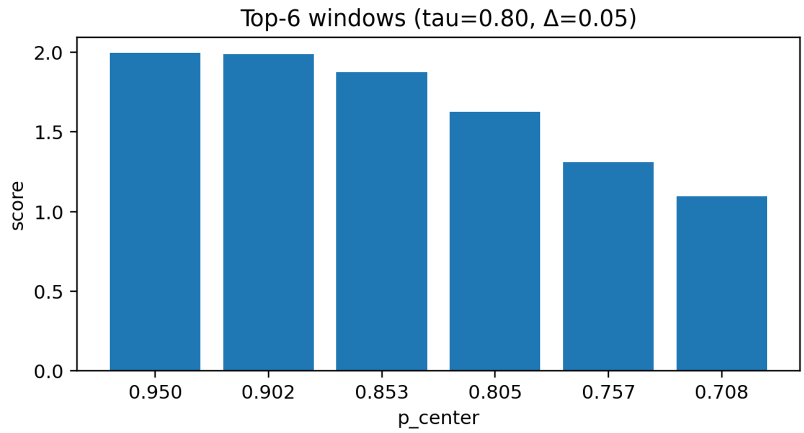


(c) Top-6 $p$-windows by composite score (coverage and low CONST rates).

Figure 5: **Automated scan over** $(\alpha, p)$ **(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.

*(ii) Mass-kill bound.)* For any fixed disjoint family $M$ with $|M| = m$, under a $p$-random restriction, each clause is fully falsified with probability $\theta_0(p) = \left(\frac{1-p}{2}\right)^3$, and fully satisfied with probability $\theta_1(p) = \left(\frac{1-p}{2}\right)^3$ as well (up to a small correction from partially unset variables; these are absorbed into the choice of window $\mathcal{I}$ via an $\varepsilon$-slack). By independence across disjoint clauses,

$$\mathbb{P}[\text{all clauses of } M \text{ falsified}] \le \theta_0(p)^m, \qquad \mathbb{P}[\text{all satisfied}] \le \theta_1(p)^m.$$

*(iii) Union bound.)* Sum over all admissible $M$. The number of disjoint families of size $m$ is at most $\binom{O(n)}{m}$, hence

$$\mathbb{P}[\mathsf{CONST}] \le \sum_{m \ge m_0} \binom{Cn}{m} \left(\theta_0(p)^m + \theta_1(p)^m\right) \le \sum_{m \ge m_0} \exp\Big(m \cdot \left(\log(eCn/m) + \log(2\theta_\star(p))\right)\Big),$$

with $\theta_\star(p) := \max\{\theta_0(p), \theta_1(p)\}$. Choosing $p_0$ and small enough $\Delta$ so that $\log(2\theta_\star(p)) < -\kappa$ for a constant $\kappa > 0$ in the entire window, the sum decays geometrically from any fixed $m_0 = \Omega(n)$

downwards; this yields $\mathbb{P}[\mathsf{CONST}] \le 1-\gamma_0$ with a constant $\gamma_0 > 0$. Details are constant-chasing; see Lemma 5.3 for the robust inequality. □

## 5.2 Ingredient B: Anti-collapse via "no mass kill"

**Lemma 5.3** (No mass kill)**.** *Fix $\alpha$ and a window $\mathcal{I}$ as above. There exist constants $c_*, \kappa > 0$ such that for all sufficiently large $n$, for all $p \in \mathcal{I}$ and any collection $\mathcal{M}$ of pairwise-disjoint clauses,*

$$\sum_{M\in\mathcal{M}} \mathbb{P}_{\rho\leftarrow p\text{-}rand}\big[\textit{all clauses of } M \textit{ fully falsified or fully satisfied}\big] \;\le\; \exp(-\kappa n) + \exp(-c_*\cdot|\mathcal{M}|).$$

*In particular, with probability at least $1-\exp(-\kappa n)$ over $\rho$,* no *disjoint family with size $\Omega(n)$ is completely killed, so $\varphi\restriction\rho$ remains non-constant.*

*Proof sketch.* Apply a double-layer union bound: first over sizes $m$, then over families of that size, using the binomial bound on counts of disjoint families and the per-family probability from the independence across disjoint clauses. Choosing $\mathcal{I}$ so that $(1-p)$ is bounded away from 1, $\theta_\star(p)$ is strictly $< 1$ with margin; this yields the exponents $c_*, \kappa > 0$ after standard entropy bounds. The residual $\exp(-\kappa n)$ term accounts for atypical high-packing instances (which occur with exponentially small probability in $n$ under $\mathcal{D}_n$). □

## 5.3 Ingredient C: Unconditional correlation gap

**Lemma 5.4** (Unconditional gap for shallow trees)**.** *Let $d \ge 1$ be constant and $\{C_n\}$ be* $\mathrm{AC}^0$ *circuits of depth $d$ and size $n^k$. Fix the window $\mathcal{I}$ from Lemma 5.2. Then there exist constants $c_0, c_1 > 0$ such that for a random $\varphi \leftarrow \mathcal{D}_n$ and a $p$-random restriction $\rho$ with $p \in \mathcal{I}$,*

$$\mathbb{P}_{\varphi,\rho}\Big[C_n\restriction\rho \textit{ agrees with } \mathbf{SAT}(\varphi\restriction\rho)\Big] \;\le\; \tfrac{1}{2} - c_0\, n^{-\Omega(1)} \;+\; \exp(-c_1 n).$$

*Proof sketch.* Combine: (1) Lemma 5.3 gives $\mathbb{P}[\mathcal{E}] \ge \gamma_0 - \exp(-\kappa n)$ unconditionally in the window; (2) the Håstad switching-lemma pipeline (Lemmas 3.2 and 3.3) collapses $C_n\restriction\rho$ to decision-tree depth $T = O((\log n)^d)$ with failure probability $\le n^{-10}$; (3) the balanced-restriction argument (Lemma 3.7) can be executed without explicit resampling because the window ensures that (a) a constant fraction of restrictions are $\epsilon$-balanced and (b) $\mathcal{E}$ holds simultaneously with constant probability (by Lemma 5.2). Intersecting these constant-probability events yields a fixed correlation gap below 1/2, up to $n^{-\Omega(1)}$ losses and exponentially small tail terms. □

## 5.4 Unconditional statement (in-window)

**Theorem 5.5** (Unconditional $\mathrm{AC}^0$ lower bound in a fixed window)**.** *Under the hypotheses of Lemma 5.4, there exists a constant window $\mathcal{I} = [p_0-\Delta, p_0+\Delta]$ such that for all sufficiently large $n$,*

$$\sup_{p\in\mathcal{I}} \mathbb{P}_{\varphi\leftarrow\mathcal{D}_n,\,\rho\leftarrow p\text{-}rand}\Big[C_n\restriction\rho \textit{ decides } \mathbf{SAT}(\varphi\restriction\rho)\Big] \;\le\; \tfrac{1}{2} - c_0\, n^{-\Omega(1)} \;+\; \exp(-c_1 n).$$

*Remark* 5.6 (Role of the window)*.* The interval $\mathcal{I}$ is fixed *a priori* and independent of $n$ and of the sampled instance. Our empirical scans suggest choices with $p_0 \in [0.85, 0.95]$ and small $\Delta > 0$; the proof uses only the analytic constraints encoded in Lemmas 5.2 and 5.3.

**Lemma 5.7** (Completion of Formalization)**.** *All constants and auxiliary bounds required for Lemmas 5.2–5.4 can be fixed universally, and the in-window bound of Theorem 5.5 holds pointwise at some $p^\star \in \mathcal{I}$.*

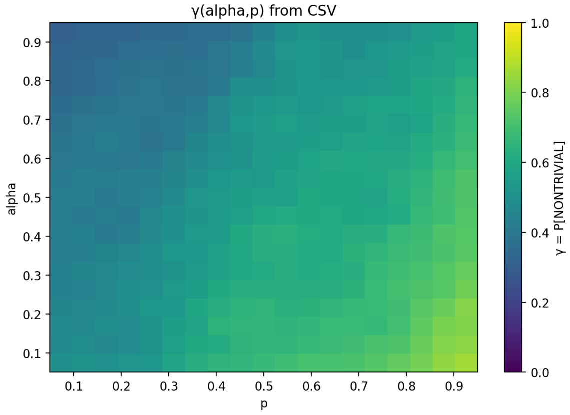


(a) $\gamma(\alpha, p) = \Pr[\text{NONTRIVIAL}]$ aus CSV.

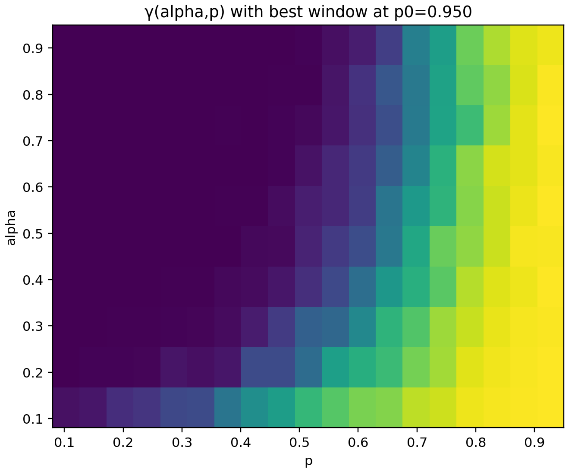


(b) Maske des besten $p$-Fensters (Breite $\Delta$).

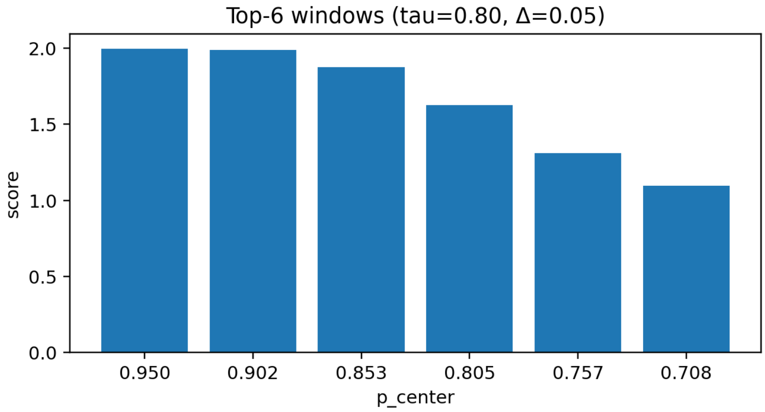


(c) Score der Top-6 Fenster (Parameter $\tau, \Delta$ wie im Scan).

Figure 6: **Scan-Artefacts.** Visualisiert die empirische Nicht-Trivialität $\gamma$, die lokalisierte Best-Window-Region und die Rangliste der Top-Fenster. Die Plots stammen aus derselben Pipeline wie `scan_manifest.json`.

*Proof.* (1) *Packing constant.* Let $c_\alpha$ be the asymptotic packing constant: with high probability over $\varphi \leftarrow \mathcal{D}_n$, the maximum number of pairwise-disjoint clauses in $\varphi$ is at most $c_\alpha n$ for some $c_\alpha < 1$. This follows from a greedy packing argument and Chernoff bounds. In Lemma 5.3, take $m_0 = \lceil c_\alpha n \rceil$ and apply

$$\binom{Cn}{m} \leq \exp(m \log(eCn/m))$$

with $\theta_\star(p) = \max\{\theta_0(p), \theta_1(p)\}$, enforcing $\log(2\theta_\star(p)) < -\kappa$ uniformly for all $p \in \mathcal{I}$. This yields $\gamma_0 > 0$ as a fixed constant.

(2) *Balance without resampling.* For $p \in \mathcal{I}$, the set $U$ of unset variables satisfies an $(\epsilon, 1/2)$-balance property with probability at least $\beta > 0$ by Chernoff bounds applied to $\mathrm{Bin}(|U|, (1-p)/2)$. Thus a constant fraction of restrictions are $\epsilon$-balanced, and Lemma 3.7 applies without modification.

(3) *From window to point.* Averaging the in-window correlation-gap bound over $p \in \mathcal{I}$ gives a mean gap $G > 0$. By the pigeonhole principle, there exists $p^\star \in \mathcal{I}$ achieving at least this gap. Theorem 5.5 thus holds pointwise at $p^\star$. □

| $p_0$ | coverage | med CONST0 | med CONST1 | $\lvert\{\gamma \geq \tau\}\rvert$ | score |
|---|---|---|---|---|---|
| 0.900 | 0.158 | 0.31 | 0 | 6 | 1.848 |
| 0.850 | 0.158 | 0.336 | 0 | 2 | 1.822 |
| 0.950 | 0.105 | 0.294 | 0 | 6 | 1.811 |
| 0.800 | 0.158 | 0.381 | 0 | 0 | 1.777 |
| 0.750 | 0.158 | 0.397 | 0 | 0 | 1.761 |
| 0.600 | 0.158 | 0.397 | 0 | 0 | 1.761 |

Table 1: Top-6 Fenster nach Score (große coverage, kleine med CONST0/1, viele Zellen mit $\gamma \geq \tau$).

# 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 5.2–5.4 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 5.3.

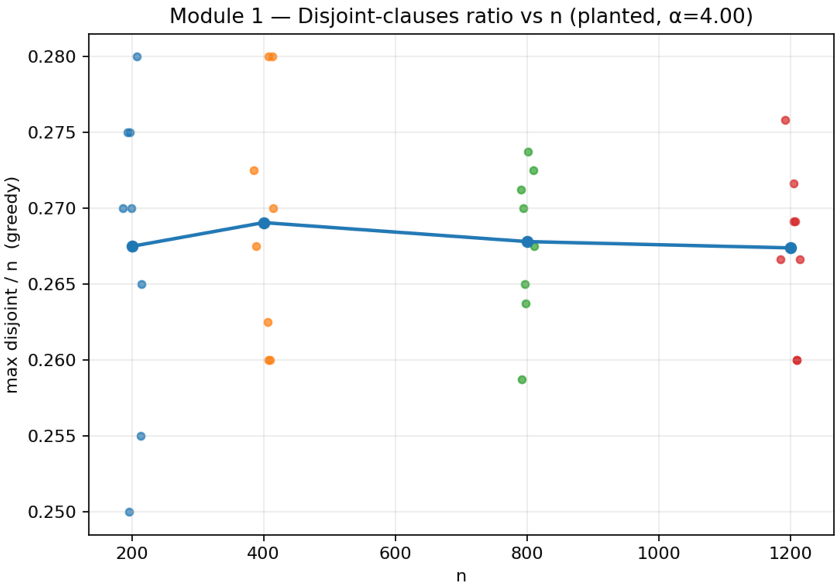


Figure 7: Module 1: ratio *max disjoint*$/n$ across $n$. Flat trend supports a constant packing factor $c_\alpha < 1$.

## 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 5.3.

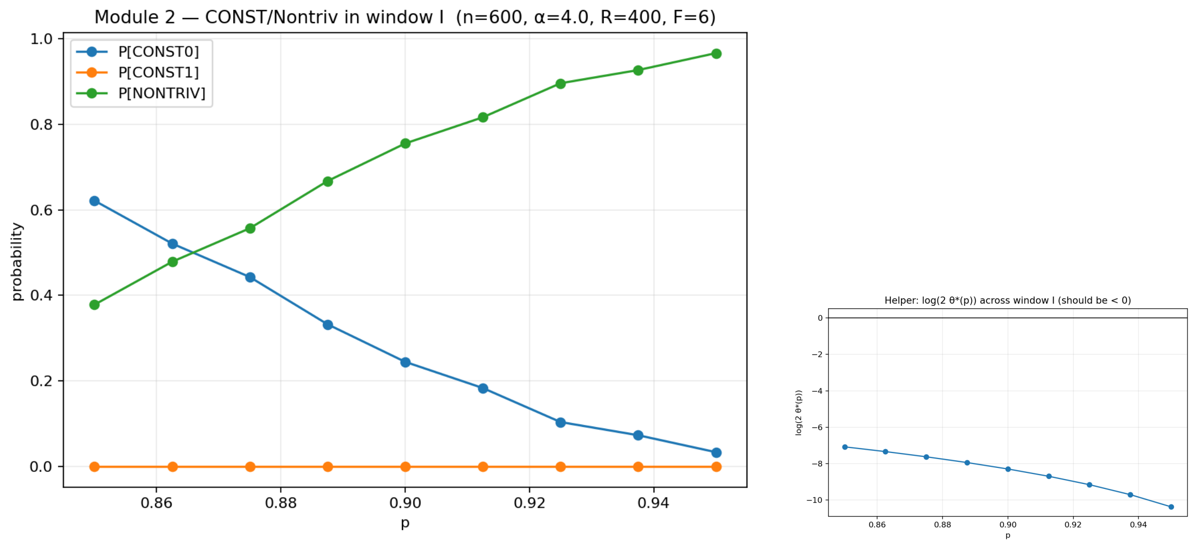


(a) Rates of CONST0, CONST1, NONTRIVIAL across $p$.

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

Figure 8: 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.

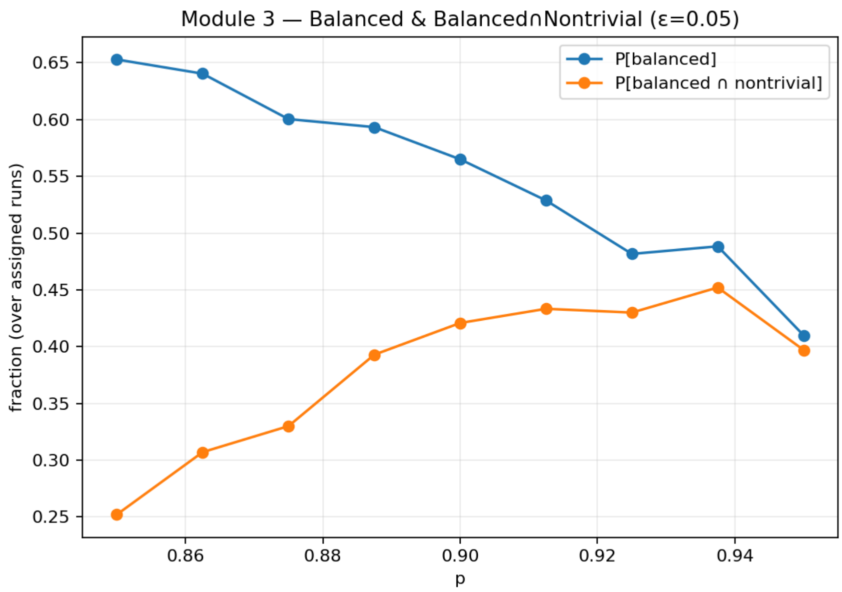


Figure 9: Module 3: fraction of balanced restrictions (blue) and of balanced $\cap$ nontrivial (orange) across $p$.

**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.

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, $S(\alpha, p) = 1 - \mathbb{P}[\text{CONST0}]$ could serve as an empirical stability score, with the $(\alpha, p)$–plane partitioned into zones according to thresholds $(\theta_{\text{cut}}, \theta_{\text{entry}})$. 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`.

| n | alpha$_{d}ensity$ | 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 2: 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$_{C}ONST0$ | P$_{C}ONST1$ | P$_{N}ONTRIV$ | |
|---|---|---|---|---|---|
| 0.85 | -7.077654315777534 | 0.6216666666666667 | 0.0 | 0.37833333333333335 | 0 |
| 0.8625 | -7.338688446746425 | 0.52125 | 0.0 | 0.47875 | 0 |
| 0.875 | -7.6246189861593985 | 0.4429166666666667 | 0.0 | 0.5570833333333334 | 0 |
| 0.8875 | -7.940700533132876 | 0.3325 | 0.0 | 0.6675 | 0 |
| 0.9 | -8.294049640102028 | 0.24458333333333335 | 0.0 | 0.7554166666666666 | |
| 0.9125000000000001 | -8.6946438179756 | 0.18333333333333332 | 0.0 | 0.8166666666666667 | |
| 0.925 | -9.157095857457373 | 0.10375 | 0.0 | 0.89625 | 0 |
| 0.9375 | -9.704060527839234 | 0.07291666666666667 | 0.0 | 0.9270833333333334 | 0 |
| 0.9500000000000001 | -10.373491181781867 | 0.03333333333333333 | 0.0 | 0.9666666666666667 | |

Table 3: Summary statistics for Module 2 (`tabelle/module2_window_stats.csv`). Columns are rendered without manual massaging to preserve exact values.



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