

PPSZ is better than you think

Dominik Scheder*

Shanghai Jiao Tong University

October 15, 2021

Abstract

PPSZ, for long time the fastest known algorithm for k-SAT, works by going through the variables of the input formula in random order; each variable is then set randomly to 0 or 1, *unless* the correct value can be inferred by an efficiently implementable rule (like small-width resolution; or being implied by a small set of clauses).

We show that PPSZ performs exponentially better than previously known, for all $k \geq 3$. For Unique-3-SAT we bound its running time by $O(1.306973^n)$, which is somewhat better than the algorithm of Hansen, Kaplan, Zamir, and Zwick.

All improvements are achieved without changing the original PPSZ. The core idea is to pretend that PPSZ does not process the variables in uniformly random order, but according to a carefully designed distribution. We write "pretend" since this can be done without any actual change to the algorithm.

Contents

1	Introduction				
	1.1	Analyzing PPSZ: permutations and forced variables	3		
	1.2	Previous improvements	4		
	1.3	Our contribution	5		
	1.4	Organization of the paper	5		
2	Brief overview of our method				
	2.1	Working with a make-belief distribution on permutations	6		
	2.2	Good make-belief distributions for PPSZ—a rough sketch	6		
3	Critical clause trees, labeled trees, and cuts				
	3.1	The Critical Clause Tree	7		
	3.2	Labeled trees and cuts	8		
4	Imp	provement for general k	10		
	4.1	The highly irregular case	10		
	4.2	Privileged variables—when $\Pr[Forced(x, \pi)]$ is already larger	12		
	4.3	The almost regular case	12		
	4.4	Using disjoint pairs to define a distribution	14		
	4.5	Two pessimistic distributions	15		

*dominik.scheder@gmail.com

5	Out	line of the case $k = 3$	18
	5.1	Canonical nodes and clause tree similarity	19
•			
6		nost regular and highly irregular formulas	20
	0.1	Intuition behind the proofs	21
7	The	e regular case: if H is large	22
	7.1	Label Density	22
	7.2	The distribution D on placements $\pi: V \to [0,1] \dots \dots \dots \dots$	23
		7.2.1 $D_{\epsilon}^{\gamma,\Box}$ on $[0,1] \times [0,1]$	23
		7.2.2 D^G : Extending $D_{\epsilon}^{\gamma,\Box}$ to paths and cycles	24
		7.2.3 The divergence from uniform to D^{G}	25
		7.2.4 Using $D^{\overline{G}}$ to sample placements π	26
	7.3	Disentangling the tree	26
		7.3.1 Making T infinite	33
		7.3.2 Dealing with bad biases: the biased nodes lemma	34
		7.3.3 When variabel(u) \in TwoCC or u has an only child	35
	7.4	General but pessimistic lower bound on $\Pr[\operatorname{Cut}(T_x)]$	35
	7.5	The case that $x \notin \text{TwoCC}$ and $\{y, z\} \in H_{\text{low}} \dots \dots \dots \dots \dots \dots$	38
	7.6	The case that $x \notin \text{TwoCC}$ and $\{y, z\} \in H_{\text{high}}$	40
	7.7	The case that $x \in \text{TwoCC}$	43
	7.8	The success probability of PPSZ under D	44
8	If tl	here are many low-degree variables	45
	8.1	$\Pr[Forced(x)]$ if $x \notin TwoCC$	46
	8.2	Summing over all $x \in V \setminus \text{TwoCC}$	53
	8.3	$\Pr[\text{Forced}(x)]$ is large when x has two critical clauses.	54
	8.4	Success Probability of PPSZ in terms of TwoCC and $ID_{0,1}$	54
A	Pro	ofs from Section 3, 4, and 6	58
В	Pro	ofs concerning the construction of D	62
С	Pro	ofs for Section 7	68
D	Pro	ofs from Section 8	80

1 Introduction

Satisfiability is a central problem in theoretical computer science. One is given a Boolean formula and asked to find a *satisfying assignment*, that is, setting the input variables to 0 and 1 to make the whole formula evaluate to 1. Or rather, determine whether such an assignment exists. A particular case of interest is CNF-SAT, when the input formula is in conjunctive normal form—that is, the formula is an AND of clauses; a clause is an OR of literals; a literal is either a variable x or its negation \bar{x} . If every clause contains at most k literals, the formula is called a k-CNF formula, and the decision problem is called k-SAT.

Among worst-case algorithms for k-SAT, two paradigms dominate: local search algorithms like Schöning's algorithm [15] and random restriction algorithms like PPZ (Paturi, Pudlák, and Zane [9]) and PPSZ (Paturi, Pudlák, Saks, and Zane [8]). Both have a string of subsequent improvements: Hofmeister, Schöning, Schuler, and Watanabe [6], Baumer and Schuler [1], and Liu [7] improve Schöning's algorithm. Hertli [5] and Hansen, Kaplan, Zamir, and Zwick [3] improve upon PPSZ.

For large k, both paradigms achieve a running time of the form $2^{n(1-c/k+o(1/k))}$, where c is specific to the algorithm (c = 1 for PPZ; $c = \log_2(e) \approx 1.44$ for Schöning; $c = \pi^2/6 \approx 1.64$ for PPSZ). Interestingly, the running time of completely different approaches like the polynomial method (Chan and Williams [2]) is also of this form. This gave rise to the Super-Strong Exponential Time Hypothesis (Vyas and Williams [16]), which conjectures that the "c/k" in the exponent is optimal; for example, it conjectures that a running time of $2^{n(1-\log(k)/k)}$ is impossible.

This paper presents an improvement of PPSZ. However, it is not an improvement of the algorithm but of its analysis. We show that the exact same algorithm performs exponentially better than previously known. Informally, PPSZ works by going through the variables in random order π ; inspecting each variable x, it tosses an unbiased coin to determine which value to assign, *unless* there is a set of at most w clauses that implies a certain value for x. Take w = 1 and this is exactly PPZ; take $w = \omega(1)$ and this is PPSZ (the exact rate by which w grows turns out to be immaterial for all currently known ways to analyze the algorithm). Our idea is to pretend that the ordering π is not chosen uniformly but from a carefully designed distribution D. This increases the success probability of PPSZ by some "bonus", which depends on D. It seems surprising that this can be done without actually changing the algorithm but it turns out to be just a straightforward manipulation, which we formally explain below in (2). There is a price to pay in terms of how much D differs from the uniform distribution: the success probability incurs a penalty of $2^{-\mathrm{KL}(D||U)}$, where $\mathrm{KL}(D||U)$ is the Kullback-Leibler divergence from the uniform distribution U to D. We focus on Unique-k-SAT, where the input formula has exactly one satisfying assignment. A "lifting theorem" by Steinberger and myself [12] shows that improving PPSZ for Unique-k-SAT automatically yields a (smaller) improvement for general k-SAT problem (without changing the algorithm).

The idea of analyzing PPSZ assuming some non-uniform distribution D on permutations and paying a price in terms of KL(D||U) is not new. It is explicit in [12] and implicit in [4] and [8]. However, all previous applications use this to deal with the case that sat(F), the set of satisfying assignments, contains multiple elements; furthermore, in [12, 4, 8], the distribution D is defined solely in terms of sat(F) and ignores the syntactic structure of F itself. In particular, in the special case that F has a unique solution, Dreverts to the uniform distribution. This paper is the first work that exploits the structure of F itself to define a distribution D on permutations, and uses this to prove a better success probability for the Unique-SAT case.

1.1 Analyzing PPSZ: permutations and forced variables

We will now formally describe the PPSZ algorithm. Let F be a formula, x a variable, and $b \in \{0, 1\}$. A formula F implies (x = b) if every satisfying assignment of F sets x to b. For example, $(x \lor y) \land (x \lor \overline{y})$ implies (x = 1) but neither (y = 0) nor (y = 1). For an integer w, we say F w-implies (x = b) if there is a set G of at most w clauses of F such that G implies (x = b).

The PPSZ algorithm with strength parameter w. Let w = w(n) be some fixed, slowly growing function. Given a CNF formula F and a permutation π , we define $ppsz(F,\pi)$ as follows: go through the variables x_1, \ldots, x_n in the order prescribed by π . In each step, when handling a variable x, check whether (x = b) is w-implied by F for some $b \in \{0, 1\}$. If so, set x to b (i.e., replace every occurrence of x in F by b). Otherwise, set x randomly to 0 or 1, with probability 1/2 each. We define ppsz(F) to first choose a uniformly random permutation π and then call $ppsz(\pi, F)$.

It should be noted that Paturi, Pudlák, Saks, and Zane in [8] formulated a stronger version of PPSZ, which tries to infer (x = b) using bounded-width resolution. The notion of *w*-implication is weaker, and a close look at their proof shows that they never use properties of resolution beyond those already possessed by *w*-implication. We assume that $\alpha = (1, ..., 1)$ is the unique satisfying assignment of *F*. This is purely for notational convenience.

Definition 1. Let π be a permutation and x a variable. Let $A \subseteq V$ be the set of variables coming before x in π , and let $F' := F|_{A\mapsto 1}$ be the restricted formula obtained from Fby setting every variable $y \in A$ to 1. If F' w-implies (x = 1) then we say x is forced under π and write Forced $(x, \pi) = 1$; otherwise we say x is guessed under π and write Forced $(x, \pi) = 0$. Let Forced $(\pi) := \sum_{x \in V} Forced(x, \pi)$.

Observation 2 ([8]). Suppose we run PPSZ with a fixed permutation π . Then ppsz(F, π) succeeds, i.e., finds α , with probability exactly $2^{-n+\operatorname{Forced}(\pi)}$.

Taking π to be a random permutation we get

$$\Pr[\operatorname{ppsz}(F) \text{ succeeds}] = \mathop{\mathbb{E}}_{\pi} \left[2^{-n + \operatorname{Forced}(\pi)} \right]$$

$$> 2^{-n + \mathop{\mathbb{E}}_{\pi}[\operatorname{Forced}(\pi)]} .$$
(1)

which follows from Jensen's inequality applied to the convex function $t \mapsto 2^t$. We are now in a much more comfortable position: $\mathbb{E}[\operatorname{Forced}(\pi)] = \sum_x \Pr[\operatorname{Forced}(x,\pi) = 1]$, and we can analyze this probability for every variable individually. Indeed, this is what Paturi, Pudlák, Saks, and Zane [8] did: they showed that $\Pr[\operatorname{Forced}(x,\pi) = 1] \geq s_k - o(1)$ if F is a k-CNF formula with exactly one satisfying assignment. Here s_k is a number defined by the following experiment: let T_{k-1}^{∞} be the complete rooted (k-1)-ary tree; pick $\pi : V(T_{k-1})^{\infty} \to [0,1]$ at random and delete every node u with $\pi(u) < \pi(\operatorname{root})$. Then s_k is the probability that the root is contained in a finite component. The o(1)term converges to 0 as w tends to infinity; thus, the growth rate of w only influences how fast this o(1) error term vanishes, but (as far as we know) does not materially influence the success probability of PPSZ. We conclude:

Theorem 3 ([8]). If F is a k-CNF formula with a unique satisfying assignment, then $\Pr[\operatorname{ppsz}(F) \text{ succeeds}]$ is at least $2^{-n+s_kn-o(n)}$. Furthermore, $s_k = \frac{\pi^2}{6k} + o(1/k)$.

1.2 Previous improvements

The analysis of Paturi, Pudlák, Saks, and Zane runs into trouble if F contains multiple satisfying assignments. In their original paper [8] they presented a workaround; unfortunately, this was somewhat technical and, for k = 3, 4, exponentially worse than the bound of Theorem 3. It was a breakthrough when Hertli [4] gave a very general analysis of PPSZ showing that the "Unique-SAT bound" also holds in the presence of multiple satisfying assignments. Curiously, his proof takes the result "Pr[Forced (x, π)] = $s_k - o(1)$ " more or less as a black box and does not ask how such a statement would have been obtained. Steinberger and myself [12] later simplified Hertli's proof and obtained a certain unique-to-general lifting theorem that is also important for our work: **Theorem 4** ([8]). If the success probability of PPSZ is at least $2^{-n+s_kn+\epsilon n}$ on k-CNF formulas with a unique satisfying assignment, for some $\epsilon > 0$, then it is at least $2^{-n+s_kn+\epsilon'n}$ on k-CNF formulas with multiple solutions, too, for some (smaller) $\epsilon' > 0$).

Concerning the Unique-SAT case, Hertli [5] designed an algorithm that is a variant of PPSZ and achieves a success probability of $2^{-n+s_3n+\epsilon n}$ for 3-CNF formulas with a unique satisfying assignment. Unfortunately, this ϵ is tiny, and his approach is extremely specific to 3-SAT, with no clear path how to generalize it to k-SAT. A result by Qin and Watanabe [11] strengthened Hertli's result a bit. More recently, Hansen, Kaplan, Zamir, and Zwick [3] published a *biased PPSZ*, a version of PPSZ in which some guessed variables are decided by a biased coin; which variables and how biased, that depends on the structure of the underlying formula. In contrast to Hertli's, their improvement is "visible": for 3-SAT, it improves the success probability from 1.3070319^{-n} from Theorem 3 to 1.306995^{-n} . Also, it works for all k (although the authors do not work out the exact magnitude of the improvement).

All this runs against the backdrop that we do not even fully understand the true success probability of PPSZ. Chen, Tang, Talebanfard, and myself [14] have shown that there are instances on which PPSZ has exponentially small success probability. Just *how* exponentially small has been tightened by Pudlák, Talebanfard, and myself [10]: we now know that PPSZ has success probability at most $2^{-(1-2/k-o(1/k))\cdot n}$ on certain instances, provided our parameter w is not too large; if you prefer PPSZ using small-width resolution, then this holds provided your width bound is really small, like $c \cdot \sqrt{\log \log n}$ [13].

1.3 Our contribution

We show that the success probability of PPSZ on k-CNF formulas is exponentially larger than 2^{-n+s_kn} . In particular,

Theorem 5 (Improvement for all k). For every $k \ge 3$ there is $\epsilon_k > 0$ such that the success probability of PPSZ on satisfiable k-CNF formulas is at least $2^{-n(1-s_k-\epsilon_k)}$.

Theorem 6 (Improved success probability for 3-SAT). The success probability of PPSZ on 3-CNF formulas with a unique satisfying assignment is at least 1.306973^{-n} .

Our improvement for Unique 3-SAT is roughly fifty percent larger than that of Hansen, Kaplan, Zamir, and Zwick [3]. This is of course nice but its importance should not be overstated. Also, for general k, it is not clear which approach gives better bounds; both approaches, in the words of [3], "only scratch the surface". Crucially, neither approach improves on the asymptotic $\frac{\pi^2}{6}$ -factor in the behavior of the savings s_k for large k.

Which approach has the greater potential? We believe our approach is *simpler* since it only focuses on the analysis and leaves the underlying algorithm unchanged. This seems like a limitation but actually gives us a certain freedom: we can exploit information gleaned from the formula, even if that information is by itself NP-hard to compute.

1.4 Organization of the paper

We outline our general idea, analyzing PPSZ under some non-uniform distribution on permutations, in Section 2. Section 3 introduces the notions of critical clause trees and "cuts" in those trees. This is mainly a review of critical clause trees as defined in [8]; however, since we will manipulate these trees extensively, we introduce more abstract and robust versions, called "labeled trees" and cuts therein. Section 4 contains our improvement for general k. With the notation introduced in Section 3, this will be rather

short. We emphasize that we strive for succinctness above all else for general k; we took no effort to optimize our improvement aimed for the simplest possible proof that *some* improvement is possible. Everything from Section 5 on deals exclusively with the case of Unique-3-SAT.

2 Brief overview of our method

2.1 Working with a make-belief distribution on permutations

Our starting point is to take a closer look a the application of Jensen's inequality:

$$\mathbb{E}_{\pi}\left[2^{-n+\operatorname{Forced}(\pi)}\right] \geq 2^{-n+\mathbb{E}_{\pi}[\operatorname{Forced}(\pi)]} .$$

This would be tight if $X := \text{Forced}(\pi)$ was the same for every permutation π . But maybe certain permutations are "better" than others. The idea is to define a new distribution D on permutations, different from the uniform distribution, under which "good" permutations have larger probability, thus $\mathbb{E}_{\pi \sim D}[X] > \mathbb{E}_{\pi \sim U}[X]$. Sadly, we have no control over the distribution of permutations: firstly, we promised not to change the algorithm; secondly, and more importantly, defining D will require some information that is itself NP-hard to come by. There is a little trick dealing with this. Generally speaking, if we want to bound the expression $\mathbb{E}_Q[2^X]$ from below but the obvious bound from Jensen's inequality, $2^{\mathbb{E}_Q[X]}$, is not good enough for our purposes, we can replace Q by our favorite P but have to pay a price. Formally:

$$\mathbb{E}_{Q}\left[2^{X}\right] = \sum_{\omega \in \Omega} Q(\omega) 2^{X(\omega)} = \sum_{\omega \in \Omega} P(\omega) \cdot \frac{Q(\omega)}{P(\omega)} 2^{X(\omega)}
= \mathbb{E}_{\omega \sim P}\left[2^{X(\omega) - \log_{2} \frac{P(\omega)}{Q(\omega)}}\right]
\geq 2^{\mathbb{E}_{P}[X] - \mathbb{E}_{\omega \sim P} \log_{2} \frac{P(\omega)}{Q(\omega)}}
= 2^{\mathbb{E}_{P}[X] - \mathrm{KL}(P||Q)}.$$
(2)

Here, $\operatorname{KL}(P||Q) := \sum_{\omega} P(\omega) \log_2\left(\frac{P(\omega)}{Q(\omega)}\right)$ is the Kullback-Leibler divergence from Q to P. If Q and P are continuous distributions (over $\Omega = [0,1]^n$, for example) with density functions f_Q and f_P , then (2) still holds, for $\operatorname{KL}(P||Q) := \int_{\Omega} f_P(\omega) \log_2\left(\frac{f_P(\omega)}{f_Q(\omega)}\right)$. This trick is not new: it plays a crucial rule in [12], and, if you look close enough, also in Hertli [4]; it appears, in simpler form, already in [8]. However, in [12, 4, 8], the distribution P is defined only to make "liquid variables" (variables x for which $F|_{x=0}$ and $F|_{x=1}$ are both satisfiable) come earlier in π and do not take the syntactic structure of F into account: they define P purely in terms of $\operatorname{sat}(F)$, the space of solutions, whereas our P will depend heavily on the structure on F as a 3-CNF formula. Our work is the first to apply this method to improving PPSZ on formulas with a unique satisfying assignment.

2.2 Good make-belief distributions for PPSZ—a rough sketch

How can we apply this idea to the analysis of PPSZ? The challenge is to find a distribution D under which $\mathbb{E}_{\pi \sim D}[\text{Forced}(\pi)]$ is larger than under the uniform distribution. Since we assume that F has the unique satisfying assignment $\alpha = (1, \ldots, 1)$, we can find, for every

variable x, a critical clause of the form $(x \vee \bar{y} \vee \bar{z})$.¹ Critical clauses play a crucial role in [8] and [3] as well. Imagine we change the distribution on permutations such that ytends to come a bit earlier than under the uniform distribution. It is easy to see that this can only decrease $\mathbb{E}[Forced(y,\pi)]$ (which is bad) and only increase $\mathbb{E}[Forced(a,\pi)]$ for all other variables. In particular, it increases $\mathbb{E}[Forced(x,\pi)]$ (which is good). Now assume the literal \bar{y} appears in a disproportionally large number of critical clauses. Then the beneficial effect of pulling y to the front of π outweighs its adverse effect. Thus, if there is a linear number of variables, each of which appears in a large number of critical clauses, then the success probability of PPSZ is larger than the baseline. This is what we call the "highly irregular case" below.

The other extreme would the "perfectly regular case", namely that every variable x has exactly one critical clause and that every negative literal \bar{y} appears in exactly two critical clauses. In this case, we find a matching M, i.e., a set of disjoint pairs of variables such that $\{y, z\} \in M$ implies that $(x \vee \bar{y} \vee \bar{z})$ is a critical clause of F, for some variable x. We then adapt the distribution on permutations such that the location of y and z is positively correlated—either they both tend to come late or they both tend to come early. This will have both (easily quantifiable) beneficial effects and (more difficult to quantify) adverse effects. However, we will see that the adverse effects can only be large if $\mathbb{E}_{\pi \sim U}[\text{Forced}(\pi)]$ is already larger than $s_k n$ under the uniform distribution.

Indeed, to obtain a stronger improvement for k = 3, we show that M need not be a matching, i.e., we allow the pairs in M to overlap. In this context, we define a certain class of distributions that might be of independent interest: for a graph G we can define a distribution on functions $\pi : V(G) \to [0, 1]$ such that $\pi(u)$ and $\pi(v)$ follow a prescribed distribution D^{\Box} on $[0, 1]^2$ whenever $\{u, v\}$ is an edge and are independent and uniform otherwise. The existence of such a distribution depends the prescribed distribution D^{\Box} and on the graph G (in particular the number of edges in G).

3 Critical clause trees, labeled trees, and cuts

3.1 The Critical Clause Tree

We assume that $\alpha = (1, \ldots, 1)$ is the unique satisfying assignment. That means that for every variable x, we can find a clause of the form $(x \lor \bar{y}_2 \lor \cdots \lor \bar{y}_k)$. This is called a *critical clause*. If there are several to pick from, we ask x to select one to be its *canonical* critical clause. For an integer $h \in \mathbb{N}$, a *critical clause tree of* x *of height* h is a rooted tree T_x of height at most h with a bunch of additional information: every node u of T_x has a *variable label* varlabel(u); if the depth of u is less than h, it has a *clause label* clauselabel(u). The tree is constructed as follows:

- Initialize T_x as consisting of a single root node, and set variabel(root) = x.
- While some node u of T_x of depth less than h does not have a clause label yet:
 - 1. Let α_u be the assignment arising from α by setting to 0 all the variables y that appear as variable labels on the path from the root to u (including both root and u). Let a := varlabel(u). In particular, $\alpha_u(a) = 0$.
 - 2. Pick a clause C that is violated by α_u (this exists since α is the unique satisfying assignment), and set clauselabel(u) := C.

¹Our informal outline assumes k = 3 to keep notation simple.

3. For each negative literal $\overline{z} \in C$, create a new child of u and give it variable label z. Note that u has at most k-1 children.

This tree is central to the analysis in [8] and also [3]. The *depth* of a node u in a tree T is the length of the path from the root to u; we abbreviate it as $d_T(u)$ or simply d(u) if T is understood. It is a (k-1)-ary tree: every node has at most k-1 children.

Observation 7. If u is a proper ancestor of v in T_x , then $varlabel(u) \neq varlabel(v)$.

In Point 2, we might have several clauses to choose from; we define the canonical critical clause tree of x of height h to be the critical clause tree T_x constructed as above, but adhering to the following tie-breaking rule in Point 2:

Canonical critical clause rule. In Point 2, if the canonical critical clause of varlabel(u) is violated by α_u , pick it as clauselabel(u); otherwise, pick the lexicographically first violated clause.

In the canonical clause tree T_x , we call a node v canonical if, for every node u on the path from the root to v (including root and v), varlabel(u) has exactly one critical clause, and this clause is clauselabel(u). Let u be a maximal non-canonical node in T_x (i.e., closest to the root), and a = varlabel(u). There are two reasons why u can be non-canonical: (1) the variable a has two or more critical clauses; (2) the clause label C := clauselabel(u) is not a critical clause (it has at least two positive literals, like ($x \lor a \lor \overline{b}$), for example). In the latter case, u has at most k - 2 children in T_x . Case (1) motivates the definition of the set TwoCC $\subseteq V$, the set of variables that have two or more critical clauses.

A path in T_x is a sequence u_0, \ldots, u_t where each u_i is the parent of u_{i+1} . That is, we never go up and then down again. By Observation 7, no variable can appear twice or more on a path in T_x . Critical clause trees are important because of the following lemma:

Lemma 8 ([8]). Suppose $w \ge (k-1)^{h+1}$, where w is the strength parameter of PPSZ. For a permutation π , let A be the set of variables coming before x in π . If every path from the root of T_x to a leaf at depth h contains a variable in A,² then Forced $(x, \pi) = 1$.

From now on, we take h = h(n) to be the largest integer such that $w \ge (k-1)^{h+1}$ and let every canonical critical clause tree be one of height h. Note that $\lim_{n\to\infty} h(n) = \infty$ because $\lim_{n\to\infty} w(n) = \infty$.

3.2 Labeled trees and cuts

We will extensively manipulate critical clause trees. Thus, it makes sense to define a generalized version:

Definition 9. A labeled tree is a rooted tree T, possibly infinite, in which

- 1. each node u has a label varlabel $(u) \in L$ in some label set $L \supseteq V$;
- 2. no label appears twice on a path; that is, if u is a proper ancestor of v in T, then $varlabel(u) \neq varlabel(v);$
- 3. each node is marked either as canonical or non-canonical; if u is non-canonical then so are all of its children;

²or rather, contains a *node* whose *variable label* is in A

4. each leaf of T is marked as either a safe leaf or an unsafe leaf.

A safe path in T is a path starting at the root that either ends at a safe leaf or is infinite. We write $\operatorname{Can}(T_x)$ to denote the set of canonical nodes in T_x . Furthermore, all labeled trees appearing in this paper are (k-1)-ary: each node has at most k-1 children.

A critical clause tree of height h becomes a labeled tree by simply marking leaves at depth h as safe leaves and all other leaves as unsafe leaves. From now on, instead of viewing π as a permutation of the variables V, we view it as a placement $\pi : V \to [0, 1]$. If π is sampled from some continuous distribution (for example the uniform distribution), then π is injective with probability 1 and defines a permutation, by sorting the variables from low- π to high- π . In fact, our π is defined on all labels, i.e., it is a function $\pi : L \to [0, 1]$.

Definition 10 (Cut and Cut_r). Let T be a labeled tree, x the label of its root, and $r \in [0,1]$. The event $\operatorname{Cut}_r(T)$ is an event in the probability space of all placements, defined as follows: mark a non-root vertex u as dead if $\pi(\operatorname{varlabel}(u)) < r$ and alive otherwise; mark root as alive. Then $\operatorname{Cut}_r(T)$ is defined to be the event that every safe path in T contains at least one dead node. $\operatorname{Cut}(T)$ is the event $\operatorname{Cut}_{\pi(x)}(T)$, i.e., all nodes u with $\pi(\operatorname{varlabel}(u)) < \pi(x)$ are marked dead.

The following observation is simply Lemma 8, framed in the new terminology:

Observation 11. Suppose $w \ge (k-1)^{h+1}$. If $\operatorname{Cut}(T_x)$ happens then $\operatorname{Forced}(x,\pi) = 1$.

For $r \in [0, 1]$, let wCut_r(T) ("weak cut") be defined as Cut_r(T), only that we additionally mark the root as dead if $\pi(x) < r$. Note that there is no corresponding event wCut(T). Weak cuts only make sense with respect to a particular $r \in [0, 1]$. Cut_r and wCut_r are intimately related: wCut_r(T) = [π (root) < r \lor Cut_r(T)]; if T_1, \ldots, T_l are the subtrees of T rooted at the children of the root, then Cut_r(T) = [wCut_r(T₁) $\land \cdots \land$ wCut_r(T_l)].

A particularly important example of a labeled tree is T_{k-1}^{∞} . This is simply an infinite complete (k-1)-ary tree: every node has k-1 children, and there are no leaves. All nodes have distinct labels. If k is understood, we simply write T^{∞} . If π is sampled uniformly at random and independently, then

$$Q_r^{(k)} := \Pr[\operatorname{Cut}_r(T_{k-1}^\infty)] = \left(\Pr[\operatorname{wCut}_r(T_{k-1}^\infty)]\right)^{k-1}$$
$$P_r^{(k)} := \Pr[\operatorname{wCut}_r(T_{k-1}^\infty)] = r \lor \Pr[\operatorname{Cut}_r(T_{k-1}^\infty)],$$

where we define $a \vee b := a + b - ab$ for $a, b \in [0, 1]$. It is a well-known result from the theory of Galton-Watson branching processes that $Q_r^{(k)}$ and $P_r^{(k)}$ are the smallest roots in [0, 1] of the equations $Q = (r + (1 - r)Q)^{k-1}$, and $P = r \vee P^{k-1}$, respectively.

Proposition 12. For $r \geq \frac{k-2}{k-1}$ it holds that $Q_r^{(k)} = P_r^{(k)} = 1$. On the interval $\left[0, \frac{k-2}{k-1}\right]$, $P_r^{(k)}$ is convex and $r \leq P_r^{(k)} \leq \frac{k-1}{k-2} \cdot r$. Also on that interval, $Q_r^{(k)} \leq \left(\frac{k-1}{k-2} \cdot r\right)^{k-1} \leq e r^{k-1}$.

For k = 3, we can give explicit solutions, on which we will heavily rely in the analysis for 3-SAT:

$$Q_r^{(3)} = \begin{cases} \left(\frac{r}{1-r}\right)^2 & \text{if } r < 1/2\\ 1 & \text{if } r \ge 1/2 \end{cases}$$

$$P_r^{(3)} = \begin{cases} \frac{r}{1-r} & \text{if } r < 1/2\\ 1 & \text{if } r \ge 1/2 \end{cases}$$

Again, if k is understood, we will simply write Q_r and P_r . Paturi, Pudlák, Saks, and Zane proved the following fact:

Lemma 13 ([8]). Let T_x be a critical clause tree of height h. Then $\Pr[\operatorname{Cut}_r(T_x)] \geq Q_r^{(k)} - \operatorname{Error}(r,h)$, for some function $\operatorname{Error}(r,h)$ that converges to 0 as $h \to \infty$, and $\Pr[\operatorname{Cut}(T_x) \geq s_k - o(1), \text{ where } s_k := \int_0^1 Q_r^{(k)} dr.$

4 Improvement for general k

4.1 The highly irregular case

We asked each variable to pick *one* of its critical clauses to be its *canonical* critical clause. This defines a directed graph on V: if the canonical critical clause of x is $(x \lor \bar{y}_1 \lor \cdots \lor \bar{y}_{k-1})$, we create arcs $(x, y_1), \ldots, (x, y_{k-1})$. This graph is the *critical clause graph*, short CCG. This graph has (k-1)n arcs, and every vertex has out-degree k-1. Let indeg(x) denote the in-degree of x in CCG. For a set $Y \subseteq V$, let e(x, Y) be the number of arcs (x, y) with $y \in Y$. Let two sets X, Y let $e(X, Y) := \sum_{x \in X} e(x, Y)$, and indeg $(Y) := \sum_{x \in V} e(x, Y) = \sum_{y \in Y} indeg(y)$. We fix some integer $k' \ge k$, to be determined later, and call a variable $x \ k'$ -heavy if indeg $(x) \ge k'$. Let HEAVY(k') be the set of all heavy variables. If k' is understood, we just write HEAVY.

Theorem 14. For every k, there is some k' and some $c_{\text{BONUSHEAVY}} > 0$ such that $\Pr[\text{PPSZ succeeds on } F] \ge 2^{-n+s_kn+c_{\text{BONUSHEAVY}} \cdot \text{indeg}(\text{HEAVY}(k')) - o(n)}$.

Proof. First, we define new distribution D on placements $\pi : V \to [0, 1]$. We fix some differentiable $\gamma : [0, 1] \to \mathbf{R}_0^+$ such that $\gamma(0) = \gamma(1) = 0$ and let $\phi := \gamma'$ be its derivative. We fix some $\epsilon > 0$ such that $1 + \epsilon \phi(r) \ge 0$ for all $r \in [0, 1]$. Let D_{ϵ}^{γ} be the distribution on [0, 1] that has density function $1 + \phi(r)$. Let D be the distribution on placements $\pi : V \to [0, 1]$ that samples $\pi(x) \in [0, 1]$ uniformly for all $x \notin \text{HEAVY}$ and $\pi(x) \sim D_{\epsilon}^{\gamma}$ for all $x \in \text{HEAVY}$. For heavy x it holds that $\Pr[\pi(x) < r] = \int_0^r (1 + \epsilon \phi(s)) \, ds = r + \epsilon \gamma(r) > r$. Loosely speaking, heavy variables x tends to come *earlier* in $\pi \sim D$ than non-heavy ones. Consequently, for heavy x, we expect $\Pr[\operatorname{Forced}(x, \pi)]$ to be *smaller* under D than under the uniform distribution:

Lemma 15. If $x \in \text{HEAVY}$ then $\Pr_{\pi \sim D}[\text{Forced}(x, \pi)]$ is at least

$$\int_0^1 Q_r(1+\epsilon\phi(r))\,dr - o(1) = s_k - \epsilon c_{\text{HEAVY}} - o(1)$$

for some c_{HEAVY} that depends only on γ and k, but not on k'.

If $x \notin \text{HEAVY}$ but has an arc (x, y) for some $y \in \text{HEAVY}$, we expect $\Pr[\text{Forced}(x, \pi)]$ to be *larger* under D, and in fact

Lemma 16. If $x \notin \text{HEAVY}$ then $\Pr_{\pi \sim D}[\text{Forced}(x, \pi)]$ is at least

 $s_k + \epsilon c_{\text{HEAVYCHILD}} e(x, \text{HEAVY}) - o(1)$

where e(x, HEAVY) is the number of arcs (x, y) with $y \in \text{HEAVY}$ and $c_{\text{HEAVYCHILD}} = \int_0^1 \gamma(r) P_r^{k-2} (1-Q_r) dr$, which only depends on γ and k but not on k'.

and

It is well-known from the theory of Galton-Watson branching processes that $Q_r < 1$ for all $r < \frac{k-2}{k-1}$ and therefore $P_r^{k-2} (1 - Q_r) > 0$ on the interval $\left[0, \frac{k-2}{k-1}\right]$. Thus, it is easy to choose some function γ and some $\epsilon > 0$ such that D_{ϵ}^{γ} is a distribution and $c_{\text{HEAVYCHILD}} > 0$. The exact shape of γ is not important in this part of the paper, where we only aim to show that *some* improvement is possible; when trying to prove a substantial improvement for k = 3, we will choose γ more carefully.

Proof of Lemma 16. Let y_1, \ldots, y_{k-1} be the labels of the children of the root of T_x and T_i be the subtree of T_x rooted at y_i . Similar to the proof of Lemma 7 in [8], we can assume that all nodes of T_x have distinct variable labels. Also, $\Pr[\pi(z) < r] \ge r$ holds for every variable z. This means we can assume (pessimistically) that $\pi(z)$ is uniform over [0, 1] for all $z \notin \{y_1, \ldots, y_{k-1}\}$. Since all labels are distinct, we have $\Pr[\operatorname{Cut}_r(T_x)] = \prod_{i=1}^{k-1} \Pr[\operatorname{wCut}_r(T_i)]$ and $\Pr[\operatorname{wCut}_r(T_{y_i})] = \Pr[\pi(y_i) < r] \lor \Pr[\operatorname{Cut}_r(T_i)]$, where we use the notation $A \lor B = A + B - AB$ for $A, B \in [0, 1]$. Furthermore, $\Pr[\pi(y_i) < r] = r + \epsilon \gamma(r)[y_i \in \operatorname{HEAVY}]$ and $\Pr[\operatorname{Cut}_r(T_i)] \ge Q_r - o(1)$. Since $(A+x) \lor B = (A \lor B) + (1-B)x$, we conclude that

$$\begin{aligned} \Pr[\operatorname{wCut}_r(T_{y_i})] &= \Pr[\pi(y_i) < r] \lor \Pr[\operatorname{Cut}_r(T_i)] \\ &\geq (r \lor Q_r) + (1 - Q_r)\epsilon\gamma(r)[y_i \in \operatorname{HEAVY}] - o(1) \\ &= P_r + \epsilon\gamma(r)(1 - Q_r)[y_i \in \operatorname{HEAVY}] - o(1) . \end{aligned}$$

Therefore,

$$\begin{aligned} \Pr[\operatorname{Cut}_r(T_x)] &\geq \prod_{y:x \to y} \left(P_r + \epsilon \gamma(r)(1 - Q_r) [y \in \operatorname{HEAVY}] \right) - o(1) \\ &\geq (P_r)^{k-1} + \sum_{y:x \to y} \epsilon \gamma(r) P_r^{k-2} (1 - Q_r) [y \in \operatorname{HEAVY}] - o(1) \\ &= Q_r + \epsilon \gamma(r) P_r^{k-2} (1 - Q_r) e(x, \operatorname{HEAVY}) - o(1) . \end{aligned}$$

Since $\pi(x)$ is uniform over [0, 1], the lemma follows from integrating the above expression over [0, 1].

We can prove Theorem 14 by summing over all variables. First, observe that $e(V \setminus \text{HEAVY}, \text{HEAVY}) = e(V, \text{HEAVY}) - e(\text{HEAVY}, \text{HEAVY}) \ge \text{indeg}(\text{HEAVY}) - e(\text{HEAVY}, V) \ge \text{indeg}(\text{HEAVY}) - k|\text{HEAVY}|$, and therefore (ignoring the o(1) term for notational convenience)

$$\begin{split} \frac{\sum_{x \in V} \Pr_{D}[\operatorname{Forced}(x,\pi)] - s_{k} n}{\epsilon} &\geq -c_{\operatorname{HEAVY}}|\operatorname{HEAVY}| + c_{\operatorname{HEAVYCHILD}} \sum_{x \notin \operatorname{HEAVY}} e(x,\operatorname{HEAVY}) - o(n) \\ &\geq -c_{\operatorname{HEAVY}}|\operatorname{HEAVY}| + c_{\operatorname{HEAVYCHILD}} e(V \setminus \operatorname{HEAVY},\operatorname{HEAVY}) \\ &\geq -c_{\operatorname{HEAVY}}|\operatorname{HEAVY}| + c_{\operatorname{HEAVYCHILD}} \left(\operatorname{indeg}(\operatorname{HEAVY}) - k|\operatorname{HEAVY}|\right) \\ &= c_{\operatorname{HEAVYCHILD}} \operatorname{indeg}(\operatorname{HEAVY}) - (c_{\operatorname{HEAVY}} + kc_{\operatorname{HEAVYCHILD}})|\operatorname{HEAVY}| \\ &\geq \left(c_{\operatorname{HEAVYCHILD}} - \frac{c_{\operatorname{HEAVY}} + kc_{\operatorname{HEAVYCHILD}}}{k'}\right) \operatorname{indeg}(\operatorname{HEAVY}) \;, \end{split}$$

where the last inequality follows from the fact that indeg(HEAVY) $\geq k'$ |HEAVY|. Since neither c_{HEAVY} nor $c_{\text{HEAVYCHILD}}$ depend on k', we can choose k' sufficiently large and make sure the expression in the parenthesis is some c > 0. This shows that $\mathbb{E}_D[\text{Forced}(\pi)] \ge s_k n - o(n) + \epsilon c \text{ indeg}(\text{HEAVY})$. Finally, using (2), we see that

 $-\log_2 \Pr[\operatorname{PPSZ succeeds}] \le n - s_k n + o(n) - \epsilon c \operatorname{indeg}(\operatorname{HEAVY}) + \operatorname{KL}(D||U)$.

Since all values $\pi(x)$ are independent under both D and U, the Kullback-Leibler divergence becomes additive, and $\operatorname{KL}(D||U) = \operatorname{KL}(D_{\epsilon}^{\gamma}||U_{[0,1]}) \cdot |\operatorname{HEAVY}|$, where $U_{[0,1]}$ is the uniform distribution on [0, 1].

Proposition 17. $\operatorname{KL}(D_{\epsilon}^{\gamma}||U_{[0,1]}) \leq \log_2(e) \epsilon^2 \Psi$ for $\Psi := \int_0^1 \phi^2(r) dr$.

Proof. We abbreviate $t := \epsilon \phi(r)$. By definition of KL for continuous distributions, we have

$$\ln(2)\operatorname{KL}(D_{\epsilon}^{\gamma}) = \int_{0}^{1} (1+t)\ln(1+t)\,dr \le \int_{0}^{1} (1+t)t\,dr = \int_{0}^{1} t\,dr + \int_{0}^{1} t^{2}\,dr = \epsilon^{2}\int_{0}^{1} \phi^{2}(r)\,dr$$

where the last equality follows from $\int_0^1 \phi(r) dr = \gamma(1) - \gamma(0) = 0$.

By choosing ϵ sufficiently small (depending only on γ and k) we can ensure that $\epsilon c \operatorname{indeg}(\operatorname{HEAVY}) - \log_2(e) \epsilon^2 \Psi = c_{\operatorname{BONUSHEAVY}} \operatorname{indeg}(\operatorname{HEAVY})$ for some $c_{\operatorname{BONUSHEAVY}} > 0$. Thus, $-\log_2 \Pr[\operatorname{PPSZ} \operatorname{succeeds}] \leq n - s_k n + o(n) - c_{\operatorname{BONUSHEAVY}} \operatorname{indeg}(\operatorname{HEAVY})$, which proves Theorem 14.

4.2 Privileged variables—when $\Pr[Forced(x, \pi)]$ is already larger

There are some abnormal cases that will interfere with our analysis below. Luckily, all those cases will imply that the variables involved already have a substantially *higher* probability of being forced.

Definition 18. A variable x is called privileged if (1) x has at least two critical clauses or it has a critical clause tree T_x such that (2) there is a variable y that appears simultaneously at depth 1 and 2 or (3) T_x has fewer than $(k-1)^2$ nodes at depth 2. Let PRIVILEGED be the set of all privileged variables.

Lemma 19. There is some $c_{\text{PRIVILEGED}} > 0$, depending only on k, such that $\Pr[\text{Forced}(x, \pi)] \ge s_k + c_{\text{PRIVILEGED}} - o(1)$ for all privileged variables x, where o(1) converges to 0 as w grows.

See Lemma A.2 in the appendix for a proof. The proof is rather straightforward. It uses some concepts and techniques we will extensively rely on in our analysis for 3-SAT. Thus, the reader who plans to venture into the 3-SAT analysis might just as well start by reading the proof of the above lemma.

4.3 The almost regular case

Theorem 14 already gives us an exponential improvement over the old analysis of PPSZ provided that indeg(HEAVY) is large (linear in n). In this section, we will come up with a corresponding bound that works well if indeg(HEAVY) is small. The final bound will then follow from a meet-in-the-middle argument.

Lemma 20. There is a collection G of canonical critical clauses such that no two clauses in G share a variable and $|G| \geq \frac{n-\text{indeg}(\text{HEAVY})}{kk'}$. Here, k' is the parameter chosen in the definition of HEAVY. *Proof.* Let C be the set of all canonical critical clauses. Greedily pick a clause $C \in C$, add it to G, and delete from C all clauses C' that share a variable with C (this obviously includes C itself). Repeat this step for as long as possible.

How many clauses does each step remove from C? Let x_1, \ldots, x_k denote the variables of C. For sure we remove the canonical critical clauses of x_1, \ldots, x_k . Additionally, we remove, for each $1 \leq i \leq k$, all canonical critical clauses containing \bar{x}_i . This removes at most a total of $k + \sum_{x \in \text{var}(C)} \text{indeg}(x)$ clauses. Thus, the total number of canonical critical clauses removed in this process is at most

$$\sum_{C \in G} \left(k + \sum_{x \in \operatorname{var}(C)} \operatorname{indeg}(x) \right)$$

$$\leq |G|k + \sum_{C \in G} \sum_{\substack{x \in \operatorname{var}(C) \\ x \notin \operatorname{HEAVY}}} (k' - 1) + \sum_{C \in G} \sum_{\substack{x \in \operatorname{var}(C) \\ x \in \operatorname{HEAVY}}} \operatorname{deg}_{\operatorname{in}}(x)$$

$$\leq |G|k + |G|k(k' - 1) + \operatorname{indeg}(\operatorname{HEAVY}) = |G|kk' + \operatorname{indeg}(\operatorname{HEAVY})$$

On the other hand, the process ends when all clauses have been removed. Since there are exactly n canonical critical clauses, it removes exactly n clauses, and therefore $|G|kk' + indeg(HEAVY) \ge n$. Solving for |G| proves the lemma.

We take a set G of canonical critical clauses as guaranteed by the lemma. We form a collection M' of disjoint pairs of variables by selecting, for each $C \in G$, two negative literals $\bar{y}, \bar{z} \in C$ and adding $\{y, z\}$ to M'.³ Call x a parent of $\{y, z\}$ if the canonical critical clause of x contains the literals \bar{y} and \bar{z} . Note that this name makes sense since in the canonical critical clause tree T_x of x, the root has two children with labels y and z, respectively. Every $\{y, z\} \in M'$ has at least one parent. We form a final collection $M \subseteq M'$ of pairs by removing each pair $\{y, z\}$ with parent x from M' if at least one of x, y, z is in PRIVILEGED. Each privileged variable x is "responsible" for the removal of at most two elements from M': one if the canonical clause of x happens to be in G; one if there is some x' with $\{x, x'\} \in M'$. Therefore,

$$|M| \ge \frac{n - \text{indeg(HEAVY)}}{kk'} - 2 |\text{PRIVILEGED}| .$$
(3)

We denote the set of all parents x of some $\{y, z\} \in M$ by PARENTM.

Theorem 21. For every c > 0 and $k \ge 3$ there is some c' > 0 such that if $|M| \ge cn$ then $\Pr[\operatorname{PPSZ succeeds}] \ge 2^{-n+s_kn+c'n-o(1)}$.

From here, the proof of Theorem 5 is simple. Let c_1 be a small constant, depending on k. If $|indeg(HEAVY)| \ge c_1 n$ then we can apply Theorem 14. If $|PRIVILEGED| \ge c_1 n$ we can apply Lemma 19. Otherwise, (3) implies that $|M| \ge c_2 n$ for some c_2 depending on c_1 and k, and $c_2 > 0$ if c_1 is small enough. We can now apply Theorem 21 and are done. This proves Theorem 5. It remains to prove Theorem 21.

Proof of Theorem 21. The upshot of Lemma 19 and Theorem 14 is that the success probability of PPSZ is exponentially larger than 2^{-n+s_kn} if at least one of |PRIVILEGED| and indeg(HEAVY) has size $\Omega(n)$. If this is not the case, then we can assume that |PRIVILEGED| is very small, and that |M| is of size $\Omega(n)$, by Lemma 20. At first reading of what follows, it might even be helpful to think of PRIVILEGED as being empty.

³For k = 3, the clause C contains exactly two negative literals; for larger k, we select two literals arbitrarily.

4.4 Using disjoint pairs to define a distribution

We choose some $\rho \leq \frac{k-2}{k-1}$, to be determined later. Define $\gamma : [0,1] \to \mathbb{R}_0^+$ by

$$\gamma(r) := egin{cases} r(
ho-r) & ext{if } r \leq
ho \ 0 & ext{if } r \geq
ho. \end{cases}$$

Let $\phi := \gamma'$ and extend this via $\phi(\rho) := -\rho$. Observe that $\phi_{\min} := \min_{r \in [0,1]} \phi(r) = -\rho$. We fix some $\epsilon > 0$. Let $D_{\epsilon}^{\gamma,\Box}$ be the distribution on $[0,1] \times [0,1]$ whose density at (r,s) is $1 + \epsilon \phi(r)\phi(s)$. This really is a density, provided that $1 + \epsilon \phi(r)\phi(s) \ge 0$ for all r and s. Let D be the distribution on placements that samples $(\pi(y), \pi(z)) \sim D_{\epsilon}^{\gamma,\Box}$ for each $\{y, z\} \in M$ and samples $\pi(x) \in [0, 1]$ uniformly for each remaining variable. All samplings are done independently. We define $\delta = \delta(r) := \epsilon |\phi_{\min}|\gamma(r) = \epsilon \rho \gamma(r)$.

Lemma 22. Let $r \in [0, 1]$. Then

$$\Pr_{D}[\operatorname{Cut}(T_{x}) \mid \pi(x) = r] \ge Q_{r} - \operatorname{DAMAGE}(r) + \operatorname{BENEFIT}(r) \cdot \mathbf{1}_{x \in \operatorname{PARENTM}} - o(1) ,$$

where

$$DAMAGE(r) := (k-1)(1-r)P_r^{k-2}\delta Q_r'$$
(4)

$$BENEFIT(r) := \epsilon \gamma^2(r)(1 - Q_r)^2 P_{r-\delta}^{k-3} .$$
(5)

and the o(1) converges to 0 as h grows.

Proof. The proof works by constructing an easy-to-analyze tree and distribution that serve as a pessimistic estimate for $\operatorname{Cut}_r(T_x)$. First, we make a simple but important observation about M and the labels in T_x :

Observation 23. Let u be a node of depth 1 in T_x and y = varlabel(u). Then $\{x, y\} \notin M$.

Proof. Suppose $\{x, y\} \in M$, for the sake of contradiction, and let a be their parent. So T_a contains two nodes X, Y at depth 1 with labels x and y, respectively. What is the clause label C_X of node X in T_a ? First, if it is C_x , the canonical critical clause of x, then X has a child with label y, since $\bar{y} \in C_x$. Second, if it is not C_x but is *some* critical clause, it must be a critical clause for x or a, meaning that x or a has at least two critical clauses. Third, if C_X is not a critical clause, then C_X has at most k-2 negative literals and X has at most k-2 children. In either case, at least one of a and x is privileged, meaning we would have eliminated $\{x, y\}$ from M. This is a contradiction.

This observation has two important consequences:

Observation 24. Suppose $x \notin \text{PARENTM}$ and let y_1, \ldots, y_{k-1} be the labels of the depth-1-nodes in T_x . Then $\pi(x), \pi(y_1), \ldots, \pi(y_{k-1})$ are independent and uniform under D.

Observation 25. Suppose $x \in \text{PARENTM}$ is a parent of $\{y, z\} \in M$. Let $y, z, v_1, \ldots, v_{k-3}$ be the labels of the depth-1-nodes in T_x . Then $\pi(x), \pi(v_1), \ldots, \pi(v_{k-3})$ are independent and uniform under D; the pair $(\pi(y), \pi(z))$ has distribution $D_{\epsilon}^{\gamma, \Box}$ and is independent from everything else.

The upshot is that we completely understand the distribution of $\pi(l)$ for the labels on the level 0 and 1 of T_x . Starting from level 2 downwards, the distribution can become complicated, so we resort to a pessimistic estimate: **Observation 26.** Let v be a variable and let $\tau : V \setminus \{v\} \to [0, 1]$ be a particular placement of all other variables. Let $r \in [0, 1]$. Then $\Pr_D[\pi(v) < r \mid \tau] \ge r - \delta$, for $\delta = \delta(r) := \epsilon |\phi_{\min}|\gamma(r)$.

Proof. If v is not contained in any pair of M, then $\Pr_D[\pi(v) < r \ | \tau] = r$. If $\{v, w\} \in M$ then $\Pr_D[\pi(v) < r \ | \tau] = \Pr_{D_{\epsilon}^{\gamma, \Box}}[\pi(v) < r \ | \ \pi(w) = \tau(w)] = r + \epsilon \gamma(r)\phi(\tau(w)) \ge r - \epsilon |\phi_{\min}|\gamma(r).$

4.5 Two pessimistic distributions

Fix $r \in [0,1]$ and let T^{∞} be the complete infinite (k-1)-ary tree in which all labels are distinct. Since all the labels are distinct, we take the liberty of writing $\pi(v)$ instead of $\pi(\text{varlabel}(v))$ for a node v in T^{∞} . We specify two distributions D_M and $D_{\overline{M}}$ on placements $L \to [0,1]$. First, we set $\Pr[\pi(v) < r] := r - \delta$ under both $D_{\overline{M}}$ and D_M for all nodes v of depth at least 2 in T^{∞} . Second, we let y_1, \ldots, y_{k-1} be the nodes of depth 1 in T^{∞} and sample all $\pi(y_i) \in [0,1]$ uniformly and independently under $D_{\overline{M}}$. Under D_M , we sample $\pi(y_3, \ldots, y_{k-1})$ uniformly and independently but sample $(\pi(y_1), \pi(y_2)) \sim D_{\epsilon}^{\gamma, \Box}$. This does not fully specify a distribution on placements $L \to [0,1]$ but it does specify the joint distribution of the events $[\pi(v) < r]$. Since we are only interested in $\Pr[\operatorname{Cut}_r(T_x)]$, this is enough. Note that we also do not need to specify a distribution for $\pi(x)$.

Observation 27. If $x \notin \text{PARENTM}$ then $\Pr[\text{Cut}(T_x) \mid \pi(x) = r] \geq \Pr_{D_{\bar{M}}}[\text{Cut}_r(T^{\infty})] - o(1)$. If $x \in \text{PARENTM}$ then $\Pr[\text{Cut}(T_x) \mid \pi(x) = r] \geq \Pr_{D_M}[\text{Cut}_r(T^{\infty})] - o(1)$.

The o(1)-term comes from the fact that T_x is a critical clause tree of the height h, whereas T^{∞} is an infinite tree.

Proposition 28. $\Pr_{D_M}[\operatorname{Cut}_r(T^\infty)] \ge \Pr_{D_{\bar{M}}}[\operatorname{Cut}_r(T^\infty)] + \operatorname{BENEFIT}(r)$ for $\operatorname{BENEFIT}(r) = \epsilon \gamma^2(r)(1-Q_r)^2 P_{r-\delta}^{k-3}$ as defined in (5).

Proof. Let T_1, \ldots, T_{k-1} be the subtrees of T^{∞} rooted at the nodes of depth 1. Let $\tau : L \setminus \{y_1, y_2\}$ be a partial placement. Observe that the distribution of τ is the same under D_M and $D_{\overline{M}}$. We call τ critical if $\operatorname{Cut}_r(T_1)$ and $\operatorname{Cut}_r(T_2)$ do not happen but $\operatorname{wCut}_r(T_3), \ldots, \operatorname{wCut}_r(T_{k-1})$ do happen. Note that this can be determined by looking at τ alone. Furthermore, $\operatorname{Pr}_{D_M}[\operatorname{Cut}_r(T^{\infty}) \mid \tau] = \operatorname{Pr}_{D_{\overline{M}}}[\operatorname{Cut}_r(T^{\infty}) \mid \tau]$ if τ is not critical. This follows from the fact that the marginal distributions of $\pi(y_1)$ and $\pi(y_2)$ are uniform under D_M . If τ is critical then

$$\Pr_{D_{\bar{M}}}[\operatorname{Cut}_{r}(T^{\infty}) \mid \tau] = \Pr_{D_{\bar{M}}}[\pi(y_{1}) < r \wedge \pi(y_{2}) < r] = r^{2} ,$$

$$\Pr_{D_{M}}[\operatorname{Cut}_{r}(T^{\infty}) \mid \tau] = \Pr_{D_{M}}[\pi(y_{1}) < r \wedge \pi(y_{2}) < r] = r^{2} + \epsilon \gamma^{2}(r) .$$

The probability that τ is critical is

$$\Pr[\neg \operatorname{Cut}_r(T_1)] \cdot \Pr[\neg \operatorname{Cut}_r(T_2)] \cdot \prod_{i=3}^{k-1} \Pr[\operatorname{wCut}_r(T_i)] \ge (1-Q_r)^2 P_{r-\delta}^{k-3},$$

and therefore

$$\Pr_{D_M}[\operatorname{Cut}_r(T^{\infty})] = \Pr_{D_{\bar{M}}}[\operatorname{Cut}_r(T^{\infty})] + \epsilon \gamma^2(r) \operatorname{Pr}[\tau \text{ is critical}] \\ \ge \Pr_{D_{\bar{M}}}[\operatorname{Cut}_r(T^{\infty})] + \epsilon \gamma^2(r)(1-Q_r)^2 P_{r-\delta}^{k-3}.$$

This completes the proof.

Proposition 29. $\Pr_{D_{\bar{M}}}[\operatorname{Cut}_r(T^{\infty})] \ge Q_r - \operatorname{DAMAGE}(r)$ for $\operatorname{DAMAGE}(r) = (k-1)(1-r)P_r^{k-2}\delta Q'_r$ as defined in (4).

Proof. If $r \geq \frac{k-2}{k-1}$ then $\delta = 0$, DAMAGE = 0, and $\Pr[\pi(v) < r] = r$ for every node v in T^{∞} . Both sides of the inequality evaluate to 1.

Otherwise, let T_1, \ldots, T_{k-1} be the subtrees of T^{∞} rooted at the depth-1-nodes of T^{∞} . Since $\Pr_{D_{\bar{M}}}[\pi(v) < r] = r - \delta$ for all nodes v of T^{∞} of depth 2 or greater, it holds that $\Pr_{D_{\bar{M}}}[\operatorname{Cut}_r(T_i)] = Q_{r-\delta}$ for $1 \leq i \leq k-1$. Since Q_r is convex on $\left[0, \frac{k-2}{k-1}\right]$, we have $Q_{r-\delta} \geq Q_r - \delta Q'_r$. Next,

$$\Pr_{D_{\bar{M}}} [\operatorname{wCut}_{r}(T_{i})] = r \vee \Pr_{D_{\bar{M}}} [\operatorname{Cut}_{r}(T_{i})] \ge r \vee (Q_{r} - \delta Q_{r}')$$
$$= r + (1 - r)(Q_{r} - \delta Q_{r}') = P_{r} - (1 - r)\delta Q_{r}'$$
$$= P_{r} \left(1 - \frac{(1 - r)\delta Q_{r}'}{P_{r}}\right) ,$$

and

$$\Pr_{D_{\bar{M}}}[\operatorname{Cut}_{r}(T^{\infty})] = \prod_{i=1}^{k-1} \Pr_{D_{\bar{M}}}[\operatorname{wCut}_{r}(T_{i})]$$

$$\geq P_{r}^{k-1} \left(1 - \frac{(1-r)\delta Q_{r}'}{P_{r}}\right)^{k-1}$$

$$\geq P_{r}^{k-1} \left(1 - \frac{(k-1)(1-r)\delta Q_{r}'}{P_{r}}\right)$$

$$= Q_{r} - (k-1)(1-r)P_{r}^{k-2}\delta Q_{r}'.$$

This completes the proof.

From here on, we estimate for $x \notin \text{PARENTM}$:

$$\Pr_{D}[\operatorname{Cut}(T_{x}) \mid \pi(x) = r] \ge \Pr_{D_{\overline{M}}}[\operatorname{Cut}_{r}(T^{\infty})] - o(1) \qquad \text{(by Observation 27)}$$
$$\ge Q_{r} - \operatorname{DAMAGE}(r) - o(1) \qquad \text{(by Proposition 29)}$$

and for $x \in \text{PARENTM}$:

$$\Pr_{D}[\operatorname{Cut}(T_{x}) \mid \pi(x) = r] \ge \Pr_{D_{M}}[\operatorname{Cut}_{r}(T^{\infty})] - o(1) \qquad \text{(by Observation 27)}$$
$$\ge \Pr_{D_{\overline{M}}}[\operatorname{Cut}_{r}(T^{\infty})] + \operatorname{BENEFIT}(r) - o(1) \qquad \text{(by Proposition 28)}$$
$$\ge Q_{r} - \operatorname{DAMAGE}(r) + \operatorname{BENEFIT}(r) - o(1) \quad .$$
$$(\text{ by Proposition 29)}$$

This concludes the proof of Lemma 22.

We obtain a lower bound on $\Pr_D[\operatorname{Cut}(T_x)]$ by integrating the bound in Lemma 22 over r:

$$\Pr[\operatorname{Cut}(T_x)] \ge s_k - \operatorname{DAMAGE} + \operatorname{BENEFIT} \cdot \mathbf{1}_{x \in \operatorname{PARENTM}} - o(1) , \qquad (6)$$

where $DAMAGE = \int_0^1 DAMAGE(r) dr$ and $BENEFIT = \int_0^1 BENEFIT(r) dr$. To simplify the integration, we will first give an upper bound on DAMAGE(r) and a lower bound on BENEFIT(r).

Proposition 30. The following bounds hold:

$$DAMAGE \le O(\epsilon \rho^{2k}) \tag{7}$$

BENEFIT
$$\geq \Omega(\epsilon \rho^{k+1})$$
, (8)

where the O hides factors depending solely on k and terms of order ρ^a for $a \ge 2k + 1$, and the Ω hides factors depending solely on k and terms of order ρ^b for $b \ge k + 2$.

Proof. We remind the reader of the definitions of DAMAGE and BENEFIT in (4) and (5):

DAMAGE
$$(r) = (k-1)(1-r)P_r^{k-2}\delta Q'_r$$

BENEFIT $(r) = \epsilon \gamma^2 (r)(1-Q_r)^2 P_{r-\delta}^{k-3}$

Both BENEFIT(r) and DAMAGE(r) vanish for $r \ge \rho$. Thus, we can replace \int_0^1 by \int_0^{ρ} . We will first bound BENEFIT(r). On the interval $[0, \rho]$, $\gamma(r) = r(\rho - r)$, and $1 - Q_r \ge 1 - Q_\rho$, and $P_{r-\delta} \ge r - \delta = r(1 - \epsilon\rho(\rho - r))$. Therefore,

$$\begin{split} \text{Benefit}(r) &\geq \epsilon r^2 (\rho - r)^2 (1 - Q_\rho)^2 r^{k-3} \left(1 - \epsilon \rho (\rho - r) \right)^{k-3} \\ &\geq \frac{1}{2} \, \epsilon r^{k-1} (\rho - r)^2 \ , \end{split}$$

for sufficiently small ϵ and ρ (smaller than a value depending solely on k). Integrating this over $r \in [0, \rho]$ shows (8).

Next, we bound DAMAGE(r) from above. It holds that $r \leq P_r \leq \frac{k-1}{k-2}r$, where the first inequality follows immediately from $P_r = r \vee Q_r$, and the second follows from the fact that P_r is convex on $\left[0, \frac{k-2}{k-1}\right]$, $P_0 = 0$, and $P_{\frac{k-2}{k-1}} = 1$. Third, we compute $Q'_r = (P_r^{k-1})' = (k-1)P_r^{k-2}P'_r \leq (k-1)\left(\frac{k-1}{k-2}\right)^{k-2}r^{k-2}P'_r$. To bound P'_r , observe again that $P'_r \leq P'_{\frac{k-2}{k-1}}$, where $P'_{\frac{k-2}{k-1}}$ is the *left derivative* since P_r is not differentiable at $r = \frac{k-2}{k-1}$. To determine the left derivative, recall that $P = P_r$ satisfies the equation

$$P = P^{k-1} + (1 - P^{k-1})r$$

and therefore

$$r = r(P) = \frac{P - P^{k-1}}{1 - P^{k-1}}$$

For $P \to 1$, the derivative of $P \mapsto r(P)$ converges to $\frac{k-2}{2(k-1)}$ (compute the derivative of r(P) and apply l'Hôpital's rule twice; then substitute P = 1). Thus, for $r \to \frac{k-2}{k-1}$, the derivative P'_r converges to the inverse thereof, to $\frac{2(k-1)}{k-2}$. Thus, $Q'_r \leq (k-1) \left(\frac{k-1}{k-2}\right)^{k-2} r^{k-2} \frac{2(k-1)}{k-2}$. Altogether,

$$\begin{aligned} \text{DAMAGE}(r) &= (k-1)(1-r)P_r^{k-2}\delta Q_r' \\ &\leq (k-1)(1-r)\left(\frac{k-1}{k-2}\right)^{k-2}r^{k-2}\epsilon\rho r(\rho-r)(k-1)\left(\frac{k-1}{k-2}\right)^{k-2}r^{k-2}\frac{2(k-1)}{k-2} \\ &= C_k\epsilon(1-r)r^{2k-3}\rho(\rho-r) \end{aligned}$$

for some constant C_k depending only on k. Integrating over $r \in [0, \rho]$ yields (7).

Combining (6) with the bounds (7) and (8) and summing over all $x \in V$, we obtain

$$\sum_{x \in V} \Pr_D[\operatorname{Cut}(T_x)] \ge s_k n - O(\epsilon \rho^{2k}) n + \Omega(\epsilon \rho^{k+1}) |M| - o(n) .$$

Finally, to bound the success probability of PPSZ using the distribution D, we need to bound $\operatorname{KL}(D||U)$ from above. By additivity of KL, we see that $\operatorname{KL}(D||U) = |M| \cdot \operatorname{KL}(D_{\epsilon}^{\gamma,\Box}||U^{\Box})$, where U^{\Box} denotes the uniform distribution on $[0,1] \times [0,1]$. The density of $D_{\epsilon}^{\gamma,\Box}$ at r,s is $1 + \epsilon \phi(r)\phi(s)$, and therefore

$$\begin{split} KL(D_{\epsilon}^{\gamma,\Box}||U^{2}) &= \frac{1}{\ln(2)} \int_{[0,1]^{2}} (1 + \epsilon \phi(r)\phi(s)) \ln(1 + \epsilon \phi(r)\phi(s)) \, ds \, dr \\ &\leq \frac{1}{\ln(2)} \int_{[0,1]^{2}} (1 + \epsilon \phi(r)\phi(s))\epsilon \phi(r)\phi(s)) \, ds \, dr \\ &= \frac{1}{\ln(2)} \int_{[0,1]^{2}} \epsilon^{2} \phi^{2}(r)\phi^{2}(s) \, ds \, dr \\ &= \frac{\epsilon^{2}}{\ln(2)} \left(\int_{0}^{1} \phi^{2}(r) \, dr \right)^{2} \\ &= \frac{\epsilon^{2} \rho^{3}}{3 \ln(2)} \, . \end{split}$$

Thus, using (2), we conclude that the success probability of PPSZ is $2^{-n+s_k n+\text{gain}}$ where

gain
$$\geq \Omega(\epsilon \rho^{k+1})|M| - O(\epsilon^2 \rho^3)|M| - O(\epsilon \rho^{2k})n$$

Choosing $\epsilon = \rho^{k-3}$, this is $\Omega(\rho^{2k-2})|M| - O(\rho^{3k-3})n$. Thus, if $|M| \ge cn$, we can choose a sufficiently small ρ , depending on k and c, and make it become at least $s_k n + c'n$, for some constant c' depending on c and k. This concludes the proof of Theorem 21.

5 Outline of the case k = 3

For k = 3 we significantly refine the above approach. The proof of Theorem 6 will be quite tedious, involving several numerical calculations. In this section, we sketch the ideas that qualitatively differ from our proof above.

First, we turn things in the "highly irregular case" somewhat upside down. Rather than pulling heavy variables towards the front of π , we will push "light" variables towards the back. A variable x is "light" if the negative literal \bar{x} occurs in fewer than two canonical critical clauses. This is more powerful but requires a series of counter-measures to balance beneficial and detrimental effects.

Second, we refine our crude notion of "privileged variables". For two variables y, z we define a quantity called LABELDENSITY (z, T_y) , which measures how often and where the variable z appears in the canonical critical clause tree T_y of y. Roughly speaking, this counts all nodes of T_y that have label z but discounts a node at depth d by a factor r^d . The idea then is if $(x \vee \overline{y} \vee \overline{z})$ is the canonical critical clause of x and LABELDENSITY (z, T_y) is large, we expect every node u of T_y with label z to correspond to a node u' of T_x of label z, and $d_{T_x}(u') = d_{T_y}(u) + 1$, since we expect T_x to contain a subtree (more or less) isomorphic to T_y . Introducing a positive correlation between $\pi(y)$ and $\pi(z)$ will increase $\Pr[\operatorname{Cut}(T_x)]$ but will decrease $\Pr[\operatorname{Cut}(T_a)]$ whenever y appears as an ancestor of z (or

vice versa) in T_a . It turns out that the damage to $\Pr[\operatorname{Cut}(T_a)]$ is roughly proportional to LABELDENSITY (z, T_y) ; if this is small, we decide to live with the damage. If it is too large, we argue that the ubiquity of label z in T_x already ensures that $\Pr[\operatorname{Cut}(T_x)]$ is larger than s_3 . This bonus is again proportional to LABELDENSITY (z, T_y) , and thus this lends itself to a meet-in-the-middle argument.

Third, and maybe most interestingly, we do not require that the pairs $\{y, z\} \in M$, on which we introduce positive correlation, be disjoint, but allow them to overlap. This means that M forms the edge set of a graph. We can introduce appropriate positive correlations as long as this graph consists of small paths and cycles. For this, we define a distribution on functions $\pi : V(G) \to [0, 1]$ such that $(\pi(u), \pi(v))$ follow some predescribed distribution $D_{\epsilon}^{\gamma, \Box}$ whenever $\{u, v\}$ is an edge of G but $\pi(u)$ is still uniform over [0, 1] even when we condition on the values of all non-neighbors of u. The existence of such a distribution is somewhat surprising but can in fact be described by a very simple formula.

5.1 Canonical nodes and clause tree similarity

Recall the definition of TwoCC, the set of variables that have two or more critical clauses. For k = 3, we will slightly generalize this definition, for technical reasons.

Definition 31. Let \tilde{F} be the CNF formula F plus all 3-clauses that can be inferred from pairs of 3-clauses of F; for example, if F contains $(x \vee \bar{y} \vee \bar{z})$ and $(a \vee \bar{x} \vee \bar{y})$, then \tilde{F} additionally contains $(a \vee \bar{y} \vee \bar{z})$. Let TwoCC be the set of variables that contain at least two critical clauses in \tilde{F} .

Note that we could easily adapt PPSZ to add all those clauses in a pre-processing step; in fact, in its original wording in [8], the algorithm *does* perform a pre-processing step, adding all clauses which can be derived by resolution of bounded width. However, since we stated PPSZ in terms of *w*-implication rather than in terms of small-width resolution (these two versions are called *weak* and *strong* PPSZ, respectively), and we promised not to change the algorithm at all, we are not allowed to add a pre-processing step. Still, note that the *w*-implication mechanism emulates it: if something is *w*-implied by \tilde{F} , it is 2w-implied by F; thus, simply increasing w by a factor of two subsumes this preprocessing step.

Next, we classify each node in a critical clause tree as *canonical* or *non-canonical*. Let u be a node in the critical clause tree T_x , and let $root(T_x) = u_0, u_1, \ldots, u_d = v$ be the path from the root to v, and let $z_i := varlabel(u_i)$, so $z_0 = x$. Furthermore, let C_i be the clause label of u_i .

Definition 32 (Canonical nodes). The node u in T_x is called a canonical node if for all $0 \le i \le d$, (1) $z_i \notin \text{TwoCC}$ and (2) the clause label of u_i is the canonical critical clause of z_i . Otherwise, we call u a non-canonical node. We denote the set of canonical nodes of T_x by $\text{Can}(T_x)$.

Note that all ancestors of a canonical node are canonical. The notion of canonical nodes is important because the canonical part of critical clause trees looks similar:

Lemma 33 (CCT similarity lemma). Let T_x be the canonical critical clause tree for T_x , u a node in T_x , v a descendant of u in T_x , and a := varlabel(u), b := varlabel(v). If v is canonical in T_x , then there is a corresponding node v' in T_a , the canonical clause tree of variable a, and the path from u to v has the same variable label sequence as the path from $\text{root}(T_a)$ to v' in T_a . In particular, varlabel(v') = b. Furthermore, v' is canonical.



CCT similarity: if the bold node of T_x with label b is canonical, then T_a contains a "copy" p' of p with the same variable and clause labels.

The lemma is restated and proved as Lemma A.1 in the appendix.

6 Almost regular and highly irregular formulas

Recall the critical clause digraph CCG defined in Section 4.1: its vertex set is V, the set of variables; for every variable x, if the canonical critical clause of x is $(x \vee \overline{y} \vee \overline{z})$, we create arcs (x, y) and (x, z). Each vertex (variable) has out-degree 2, giving a total of 2n arcs. It also has *average* in-degree 2, although some variables might have in-degree 0, 1, or 3 or more. For $i \in \mathbb{N}_0$, let ID_i be the set of variables x with $\deg_{in}(x) = i$. Let $ID_{0,1} = ID_0 \cup ID_1$.

We define the sibling graph SG = (V, E), an undirected multigraph on the set of variables V: for every $x \in V$, let $(x \vee \overline{y} \vee \overline{z})$ be its canonical critical clause. We add the edge $\{y, z\}$ to E. Note that |E| = n (counting parallel edges by their multiplicity). What is deg_G(y)? It is the number of variables x in whose canonical critical clause y appears; thus, it is deg_{in}(y), its in-degree in the (directed) critical clause graph. The next lemma is a more precise version of Lemma 20.

Lemma 34. There is a set $H \subseteq E(SG)$ of maximum degree 2 (i.e., H consists of paths and cycles) with $|H| \ge n - |ID_1| - 2 |ID_0|$.

See Lemma A.3 in the appendix for a proof. From here on, our path is as follows: we distinguish between the "almost regular" case, when $|\text{ID}_1| + 2 |\text{ID}_0|$ is small, the sibling graph SG = (V, E) is "almost 2-regular", and we can find a large maximum-degree-2 subgraph (V, H); and the "highly irregular" case, when $|\text{ID}_1| + 2 |\text{ID}_0|$ is large. In either case, we define an appropriate make-believe distribution D and show that PPSZ outperforms its earlier analysis. The particular definition of D relies on the set $H \subseteq E(SG)$ in the almost regular case and on the sets ID_0 and ID_1 in the highly irregular case. Note that $H, \text{ID}_0, \text{ID}_1$ are by themselves hard to compute.⁴

Theorem 35 (PPSZ on almost regular formulas). Let $H \subseteq E(SG)$ be a subset of edges of the sibling graph such that (V, H) has maximum degree 2. Then the success probability of PPSZ is at least $2^{-n+s_3n+gain_1-o(n)}$ for

$$\operatorname{gain}_1 \ge \frac{|H|}{10118} - \frac{n}{41391}$$
.

⁴To be honest, we have not formally proved hardness; but it seems plausible that even deciding whether a particular clause is a critical clause is coNP-hard.

Theorem 36 (PPSZ on highly irregular formulas). The success probability of PPSZ is at least $2^{-n+s_3n+gain_2-o(n)}$

$$\operatorname{gain}_2 \ge \frac{|\mathrm{ID}_1| + 2 |\mathrm{ID}_0|}{1380}$$

where ID_i is the set of variables with in-degree *i* in the critical clause graph.

We define irr := $\frac{|\text{ID}_1|+2|\text{ID}_0|}{n}$ as a measure of how irregular the critical clause graph is. Combining the two theorems, we see that the success probability of PPSZ is at least $2^{-n+s_3n+\text{gain}-o(n)}$, for

$$gain := n \cdot \max\left(\frac{|H|/n}{10118} - \frac{1}{41391}, \frac{\text{irr}}{1380}\right)$$

$$\geq n \cdot \max\left(\frac{1 - \text{irr}}{10118} - \frac{1}{41391}, \frac{\text{irr}}{1380}\right) \qquad \text{(by Lemma 34)}$$

$$\geq \frac{n}{15218} .$$

Thus, the success probability of PPSZ is at least $\Omega(1.306973^{-n})$, which proves Theorem 6.

6.1 Intuition behind the proofs

The proofs of Theorem 35 and Theorem 36 are both rather technical, involving lots of calculations, but rest on a handful of simple ideas. To understand the idea behind the proof of Theorem 36, suppose we sample $\pi(x)$ not uniformly but from some distribution D under which $\pi(x)$ tends to be larger; that is, x tends to come *later* in the permutation. This increases $\Pr[\operatorname{Cut}(T_x)]$ but decreases $\Pr[\operatorname{Cut}(T_a)]$ for all critical clause trees T_a where x appears as a non-root label. If $x \in \operatorname{ID}_{0,1}$, there should not be too many of those: after all, there is at most one a with $(a, x) \in E(CCG)$. However, although we know that x has at most one "parent" in CCG, we have no control over the number of "grandparents", i.e., variables b with $(b, a), (a, x) \in E(CCG)$. To remedy the detrimental effect on T_b , we make sure that $\pi(a)$ tends to be a bit smaller than under the uniform distribution. Overall, we choose these biases such that the effect on T_x is beneficial (it increases $\Pr[\operatorname{Cut}(T_x)]$), a bit detrimental effect on T_a . And since there is at most one such a (since $x \in \operatorname{ID}_{0,1}$), the overall net effect is beneficial.

The idea behind the proof of Theorem 35 is more complex. We choose the distribution D such that, whenever $\{y, z\} \in H$, the pair $(\pi(y), \pi(z))$ are positively correlated but both marginal distributions are uniform over [0, 1]. If $x \vee \overline{y} \vee \overline{z}$ is the canonical critical clause of x, this turns out to increase $\Pr[\operatorname{Cut}(T_x)]$ by some amount. Unfortunately, if y appears "above" z in some other critical clause tree T_a , the effect is detrimental, i.e., it decreases $\Pr[\operatorname{Cut}(T_a)]$. We have to resort to a meet-in-the-middle argument: if T_a includes only few positively correlated ancestor-descendant pairs with labels y, z, the detrimental effect will be small; if there are many such pairs in T_a , then T_x will contain many, too, meaning the label z appears often in T_x . We can show that this has a beneficial effect on T_x . Of course, we have to add those effects for all pairs $\{y, z\} \in H$. We are lucky: the detrimental effects turn out to be roughly additive, while the beneficial effects are roughly super-additive.

The next two sections are devoted to prove Theorem 35 (Section 7) and Theorem 36 (Section 8). We should note that Theorem 35 is much more difficult to prove; however,

we feel that the "almost regular case" is the more interesting one, also closer to what we suspect are worst-case instances. Therefore, we decide to prove Theorem 35 first.

7 The regular case: if H is large

This section contains the proof of Theorem 35. Recall the set H in the theorem, a set of edges in the sibling graph with that (V, H) has maximum degree 2. Also, recall the definition of TwoCC $\subseteq V$, the set of variables that have two or more critical clauses. We say an edge $\{y, z\}$ in the sibling graph is TwoCC-free if $y \notin$ TwoCC and $z \notin$ TwoCC and, for all $x \in V$ whose canonical critical clause is $(x \lor \overline{y} \lor \overline{z})$, also $x \notin$ TwoCC. Let $H_{\text{free}} \subseteq H$ be the set of all TwoCC-free edges in H.

Observation 37. $|H_{\text{free}}| \ge |H| - 3 |\text{TwoCC}|$.

Indeed, let $x \in \text{TwoCC}$ and $(x \lor \overline{y} \lor \overline{z})$ be its critical clause. The variable x is responsible for at most three edges of H becoming "un-free": $\{y, z\}$ (if this happens to be in H), and $\{x, x'\}$ for each of the at most two neighbors x' of x in H.

7.1 Label Density

Let y, z be variables and T_y the critical clause tree of y. We need some way to quantitatively measure how prominently z features in T_y . To this end, we define the *label density* of z in T_y , denoted LABELDENSITY (z, T_y) , by

$$LABELDENSITY(z, T_y, r) := \sum_{\substack{v \in \operatorname{Can}(T_y) \\ \operatorname{varlabel}(v) = z}} \frac{(1-2r)^2}{(1-r)^3} \cdot r^{d(v)+1}$$
(9)

LABELDENSITY
$$(z, T_y) := \int_0^{1/2} \text{LABELDENSITY}(z, T_y, r) dr$$
, (10)

where $d(v) := d_{T_y}(v)$ is the depth of v in T_y (i.e., the distance from v to the root of T_y). Note that the sum in (9) goes over all canonical nodes of T_y whose variable label is z. We choose some threshold THR > 0. Our final choice will be THR := $\frac{2}{0.9 \cdot 10118} \approx \frac{1}{4553}$. Using THR and the notion of label density, we define disjoint subsets of H_{free} called H_{high} , H_{low} , and H_{rest} : For each edge $\{y, z\} \in H_{\text{free}}$ such that LABELDENSITY $(z, T_y) \geq$ THR or LABELDENSITY $(y, T_z) \geq$ THR, insert $\{y, z\}$ into H_{high} . Next, for each $\{y, z\} \in H_{\text{high}}$, assume without loss of generality that LABELDENSITY $(z, T_y) \geq$ THR; note that there is at most one other edge $\{z, z'\} \in H_{\text{free}}$; if there is one, add this edge $\{z, z'\}$ to H_{rest} . The set $H_{\text{free}} \setminus H_{\text{high}} \setminus H_{\text{rest}}$ consists of cycles and paths and has at least $|H_{\text{free}}| - 2|H_{\text{high}}|$ edges, since $|H_{\text{rest}}| \leq |H_{\text{high}}|$; we can remove a $\frac{1}{18}$ -fraction of $H_{\text{free}} \setminus H_{\text{high}} \setminus H_{\text{rest}}$ to make sure that the remaining edges form connected components with at most 17 edges each. Let H_{low} be the set of remaining edges, and observe that

$$\frac{18}{17} |H_{\text{low}}| + 2 |H_{\text{high}}| + 3 |\text{TwoCC}| \ge |H|$$
(11)

Informal synopsis of what follows. We want to define a distribution D on placements $\pi : V \to [0,1]$ under which $\pi(y)$ and $\pi(z)$ are positively correlated if $\{y,z\} \in H_{\text{low}}$ and independent otherwise. How will this affect the probabilities $\Pr[\text{Cut}(T_a)]$? As a yardstick, the positive correlation will *boost* $\Pr[\text{Cut}(T_a)]$ if y, z appear as siblings in T_a , i.e., children

of the same parent, and in particular if y, z are the children of the root. Thus, it will in particular boost $\Pr[\operatorname{Cut}(T_x)]$ when $(x \vee \overline{y} \vee \overline{z})$ is the canonical critical clause of x. It will be straightforward to quantify this boost.

The are two challenges to this approach. First, the positive correlation between $\pi(y)$ and $\pi(z)$ will decrease $\Pr[\operatorname{Cut}(T_a)]$ if y is an ancestor of z in T_a (to be more precise, if there are nodes u, v in T_a with varlabel(u) = y and varlabel(v) = z), or vice versa. The magnitude of this detrimental effect will be stronger the closer v is to the root of T_a . This points to a way to bound this effect: if u has not one but many descendants with label z, then we will see that the label z appears very frequently in T_y , too (formalized in the CCT similarity lemma, Lemma 33; this in turn means that LABELDENSITY (z, T_y) is large; if it is greater than THR then $\{y, z\}$ is in H_{high} , so $\pi(y), \pi(z)$ are independent. In this case, we will show that $\Pr[\operatorname{Cut}(T_x)]$ gets a direct boost from z appearing very often in T_x .

The second problem is that we want $\pi(y), \pi(z)$ to be independent if $\{y, z\} \notin H_{\text{low}}$. The naive approach would be to make sure H_{low} is a set of paths and then, on each H_{low} -path x_0, \ldots, x_t , make the $\pi(x_i)$ form a Markov chain such that $\pi(x_{i-1})$ and $\pi(x_i)$ are positively correlated in the way we want. This approach has two drawbacks; first, $\pi(y)$ and $\pi(z)$ will be dependent whenever y, z lie on the same H-path; second, if the original H_{low} is a set of 3-cycles, for example, we would lose 1/3 of all edges to make sure it becomes a set of paths. This approach is possible but makes us lose quite a bit of oomph. Instead, we define a "not-quite Markov chain"; given a distribution D_2 on $[0,1] \times [0,1]$ and a graph G, we define a distribution D_G on functions $\pi : V(G) \to [0,1]$ such that $\pi(u), \pi(v)$ follow D_2 if $\{u, v\} \in E$ and are independent and uniform if not. The distribution D_G will possess all the independence properties we need. However, it works only if the number of edges in G is small enough. For our choice of D_2 , "small enough" means "at most 17". This is why we make sure that every connected component of H_{low} has at most 17 edges.

For the rest of this section, we will do five things: in Section 7.2 we define the distribution D; in Sections 7.3 and 7.4 we bound the detrimental effects should z be a descendant of y in some other critical clause tree T_a ; Section 7.5 quantifies how much the positive correlation of $\pi(y)$ and $\pi(z)$ boosts $\Pr[\operatorname{Cut}(T_x)]$; Section 7.6 quantifies the boost to $\Pr[\operatorname{Cut}(T_x)]$ if LABELDENSITY (z, T_y) is large; Section 7.7 analyzes the cut probability of T_x if $x \in \operatorname{TwoCC}$; finally, in Section 7.8, we put the bounds of the four earlier sections together and prove Theorem 35.

7.2 The distribution D on placements $\pi: V \to [0,1]$

Let $\gamma : [0,1] \to \mathbb{R}_0^+$ be a continuous function with $\gamma(0) = \gamma(1) = 0$. Let $\phi := \gamma'$ be its derivative. We are fine with γ failing to be differentiable, as long as this happens for only a constant number of points. For example, γ may be piecewise linear.

7.2.1
$$D_{\epsilon}^{\gamma,\Box}$$
 on $[0,1] \times [0,1]$

Definition 38. For $\epsilon \in \mathbb{R}$, let D_{ϵ}^{γ} be the distribution on [0,1] with probability density $1 + \epsilon \phi(r)$ and cumulative probability distribution $\Pr_{X \sim D_{\epsilon}^{\gamma}}[X < r] = r + \epsilon \gamma(r)$. Let $D_{\epsilon}^{\gamma, \square}$ be the distribution on $[0,1] \times [0,1]$ whose density at (x,y) is $1 + \epsilon \phi(x)\phi(y)$.

Note that $\int_0^1 \phi(x) dx = \gamma(1) = 0$ and therefore $1 + \epsilon \phi(x)$ and $1 + \epsilon \phi(x) \phi(y)$ are indeed probability densities, provided that they are non-negative for all values x, y.

Proposition 39. Suppose $(X, Y) \sim D_{\epsilon}^{\gamma, \Box}$ and $r \in [0, 1]$. Then

- 1. $\Pr[X < r] = r$ and similarly $\Pr[Y < r] = r$.
- 2. $\Pr[X, Y < r] = r^2 + \epsilon \gamma^2(r).$
- 3. $\Pr[X < r \mid Y = b] = r + \epsilon \phi(b) \gamma(r)$.
- 4. $\Pr[X < r \mid Y \ge r] = r \frac{\epsilon \gamma^2(r)}{1-r}.$

The proof of this proposition is straightforward. In our application, we choose

$$\gamma(r) := \begin{cases} r(1-2r)^{3/2} & \text{if } r \le 1/2, \\ 0 & \text{if } r > 1/2. \end{cases}$$
(12)

Note that $\gamma(r)$ is continuous differentiable, and $\phi(r) := \gamma'(r)$ is continuous. One checks that $-1/\sqrt{5} \leq \phi(r) \leq 1$ for all $r \in [0,1]$. We will also choose some $\epsilon \leq 0.13$, and thus $1 + \epsilon \phi(x)$ and $1 + \epsilon \phi(x)\phi(y)$ are really probability density functions on [0,1] and $[0,1] \times [0,1]$. Still, we will try to keep definitions and statements as general as possible; if something only holds for this particular choice of $\gamma(r)$, we will explicitly say so.

7.2.2 D^G : Extending $D_{\epsilon}^{\gamma,\Box}$ to paths and cycles

Our goal is to define a distribution D on placements $\pi : V(H_{\text{low}}) \to [0, 1]$ such that $(\pi(y), \pi(z)) \sim D_{\epsilon}^{\gamma, \Box}$ whenever $\{y, z\} \in V(H_{\text{low}})$, and "somewhat independent" otherwise. The obvious way to do so would be to make H_{low} a collection of paths (by removing one edge per cycle) and, for each path u_0, \ldots, u_t , define a Markov chain by sampling $x_0 \in [0, 1]$ uniformly at random and x_i from [0, 1] with density $f(x_i) := 1 + \epsilon \phi(x_{i-1})\phi(x_i)$, then setting $\pi(u_i) := x_i$. This has two drawbacks: first of all, it is not clear how to build a Markov chain on a *cycle*; second, non-neighbors like x_i and x_{i-2} will not be independent; these "second-degree dependencies" can in theory be dealt with but turn out to be very nasty in practice.

Instead, we define something that could be described as a "first-degree approximation" of a Markov chain, and which generalizes nicely to cycles and, in fact, general graphs. For a graph G = (V, E), we define a distribution D^G for $\mathbf{X} = (X_v)_{v \in V}$ ranging over $[0, 1]^V$ with density function

$$f_{D^G}(\mathbf{x}) = 1 + \epsilon \sum_{\{u,v\} \in E(G)} \phi(x_u) \phi(x_v) \ .$$

This is a density function provided that it is non-negative everywhere. For our particular choices $\gamma(r)$ as defined in (12) and $\epsilon \leq 0.13$, one checks that this holds provided that $|E| \leq 17$. For each connected component C of H_{low} (which is either a path or a cycle and has at most 17 edges) we sample π independently with density f_{DC} . The following theorem lets us compute conditional probabilities under D^G :

Theorem 40. Let $A_v \subseteq [0,1]$ for $v \in V(G)$ be non-empty intervals; let $V(G) = K \uplus I$ and $E_I := \{\{u, v\} \in E \mid u, v \in I\}$. Then

$$\Pr[X_k \in A_k \ \forall k \in K \mid X_i \in A_i \ \forall i \in I] = \prod_{k \in K} \mu(A_k) \cdot \left(1 + \epsilon \frac{\sum_{\{u,v\} \in E \setminus E_I} T_u T_v}{1 + \epsilon \sum_{\{u,v\} \in E_I} T_u T_v}\right) ,$$

where μ is the Lebesgue measure on [0,1] and $T_u := \mathbb{E}_{x \in A_u}[\phi(x)]$, the expectation being taken with respect to the uniform distribution on A_u .

See Theorem B.1 in the appendix for a proof. Intuitively, the product $\prod_{k \in K} \mu(A_k)$ is what the probability were under the uniform distribution, i.e., when $\epsilon = 0$; pretending that the denominator in the fraction above is 1, the theorem basically states that the probability of $\prod A_k$ is influenced only by edges touching K.

Corollary 41. The marginal distribution of X_u under D^G is the uniform distribution; if $\{u, v\} \in E(G)$, then the marginal distribution of (X_u, X_v) under D^G is $D_{\epsilon}^{\gamma, \Box}$.

Proof. The first part follows from the second, by an additional marginalization step; to see why the second part is true, we compute the marginal distribution of (X_u, X_v) by setting $I = V \setminus \{u, v\}$, $A_w := [0, 1]$ for all $w \in I$, and A_u, A_v as we like; we apply the theorem and observe that the $w \in I$ could just as well not be there: it is the same under D^G as it would be if G consisted of a single edge $\{u, v\}$, i.e., under $D_{\epsilon}^{\gamma, \Box}$. Since the distribution of (X_u, X_v) is characterized by the probability of sets $A_u \times A_v$ for A_u, A_v being intervals,⁵ this concludes the proof.

The following corollary describes an important special case and is proved as Corollary B.2 in the appendix:

Corollary 42. Let $u \in V(G)$ and $A_u := [0, r]$; for each other $v \in V \setminus \{u\}$, suppose A_v is one of $\{r\}$, [0, r], [r, 1], or [0, 1]. Define $T_v^- = \min(0, T_v)$. Then

$$\Pr[X_u \in A_u \mid X_v \in A_v \text{ for all other } v] \ge r + \frac{\epsilon \gamma(r) \sum_{v:\{u,v\} \in E} T_v^-}{1 - \frac{2}{25} \epsilon(|E| - 1)}$$

holds for our particular choice $\gamma(r) = r(1-2r)^{3/2}$. Furthermore, if $|E(G)| \leq 17$ and $\epsilon \leq 0.13$, this is at least

$$r+1.2\,\epsilon\gamma(r)\sum_{v:\{u,v\}\in E}T_v^-$$

7.2.3 The divergence from uniform to D^G

For this section, define $m_d := \mathbb{E}[\phi^d(x)]$. This is the d^{th} moment of ϕ if we view $\phi(X)$ as a random variable with X uniform over [0,1]. Note that $m_1 = 0$. We write $f_{\text{KL}}(\epsilon) := (1-\epsilon)\ln(1-\epsilon) + \epsilon$.

Lemma 43. If $|\phi(x)| \leq 1$ for all $x \in [0,1]$ then $\operatorname{KL}(D_{\epsilon}^{\gamma}||U) \leq \frac{m_2}{\ln(2)} \cdot f_{\operatorname{KL}}(\epsilon)$. Using $\ln(1-\epsilon) \leq -\epsilon - \epsilon^2/2$, this is at most to $\frac{m_2}{2\ln(2)}(\epsilon^2 + \epsilon^3)$.

See Lemma B.3 for a proof. Next, consider the distribution D^G described above. We want to bound KL(D||U) in terms of ϵ and t.

Lemma 44. Let G be a cycle or a path, consisting of at most t edges. For $\gamma(r) = r(1-2r)^{3/2}$, $\phi(r) = \gamma'(r)$, $\epsilon \leq 0.13$, and $t \leq 17$, it holds that $\text{KL}(D^G||U) \leq 0.0064 \epsilon^2 t$.

For the curious reader who wants to experiment with different choices for $\gamma(r)$, here is a more general but non-rigorous version. Let $m_2 := \mathbb{E}_{x \in [0,1]}[\phi^2(x)]$. Then $\mathrm{KL}(D||U)$ should be bounded by $\frac{\epsilon^2 t m_2^2}{2 \ln(2)}$ times something only a bit larger than 1. The "only a bit" depends on ϵ , t, and the choice of γ . See Lemma B.4 for a proof.

⁵Technically, we should specify that D^G is a probability distribution on $[0, 1]^V$ with the σ -algebra being the Borel sets; in practice, all sets we are interested in are intervals, anyway.

7.2.4 Using D^G to sample placements π

As stated before, we start sampling $\pi \sim D$ by sampling it independently on every connected component of H_{low} . Next, for all $x \in V \setminus \text{TwoCC} \setminus V(H_{\text{low}})$, sample $\pi(x) \in [0, 1]$ uniformly and independently. It remains to sample $\pi(x)$ for $x \in \text{TwoCC}$. We define

$$\gamma_{\text{TwoCC}}(r) := \max\left(0, 40r^{7/2}(1-2r)^2\right)$$

and sample $\pi(x) \sim D_{\epsilon}^{\gamma_{\text{TwoCC}}}$ independently for each $x \in \text{TwoCC}$. This completes the description of D.

Why this particular choice? Why do we choose $\gamma(r) = r(1-2r)^{3/2}$? For one, note that $Q_r = \Pr[\operatorname{Cut}_r(T^{\infty})]$ becomes 1 for $r \geq 1/2$; thus, there is no point in distorting the distribution of $\pi(x)$ in the range [1/2, 1]. Second, we need $\gamma(0) = \gamma(1) = 0$ otherwise $1 + \epsilon \phi(x)\phi(y)$ is not a density function on $[0, 1] \times [0, 1]$. So it's a natural choice to make $\gamma(r) = 0$ for all $r \in [1/2, 1]$ and thus keep $\pi(x)$ uniform on [1/2, 1]. Third, and least obvious, bounding the detrimental effects that a positive correlation between $\pi(y)$ and $\pi(z)$ has on $\Pr[\operatorname{Cut}_r(T_a)]$ becomes harder as r grows; for our method of bounding those effects, we need that $\phi(r) = \gamma'(r)$ vanishes at r = 1/2. The factor $(1 - 2r)^{3/2}$ turns out to work. In fact, the 3/2 in the exponent could be replaced by a different constant, but 1 or less would not allow us to bound the detrimental effects, and too large a constant would reduce the beneficial effects. Other than that, we have not checked whether the constant 3/2 is optimal, neither whether the general shape of $\gamma(r)$ can be improved.

As for the weird choice $\gamma_{\text{TwoCC}}(r) = 40r^{7/2}(1-2r)^2$, it will simply turn out to be good enough; that is, with that choice, variables in TwoCC do not make any trouble, and the worst case of the analysis happens for TwoCC = \emptyset . Tinkering with γ might very well improve the eventual running time result; tinkering with γ_{TwoCC} most likely will not. Next, we need to bound KL(D||U).

Lemma 45. For our particular choice of γ and γ_{TwoCC} , $\epsilon \leq 0.13$, and every component of H_{low} having at most 17 edges, the Kullback-Leibler divergence KL(D||U) from U to D is at most

$$\operatorname{KL}(D||U) \le 0.0064 \epsilon^2 |H_{\text{low}}| + \frac{5}{48 \ln(2)} f_{\text{KL}}(\epsilon) |\operatorname{TwoCC}|$$

Proof. D is independent on every connected component of H_{low} and on every variable outside $V(H_{\text{low}})$. Since KL is additive for independent random variables, we get

$$\begin{split} \mathrm{KL}(D||U) &= \sum_{x \in \mathrm{TwoCC}} \mathrm{KL}(D_{\epsilon}^{\gamma_{\mathrm{TwoCC}}}||U) + \sum_{\substack{C: \mathrm{connected}\\\mathrm{component} \text{ of } H_{\mathrm{low}}}} \mathrm{KL}(D^{C}||U) \\ &\leq \frac{m_{2}(\phi_{\mathrm{TwoCC}})}{\ln(2)} \cdot f_{\mathrm{KL}}(\epsilon) |\mathrm{TwoCC}| + 0.0064 \, \epsilon^{2} \sum_{\substack{C: \mathrm{connected}\\\mathrm{component} \text{ of } H_{\mathrm{low}}}} |E(C)| \ , \end{split}$$

by Lemma 43 and 44, where $m_2(\phi_{\text{TwoCC}}) := \int_0^1 \phi_{\text{TwoCC}}^2(r) dr = \frac{5}{48}$ for our choice $\gamma_{\text{TwoCC}}(r) = 40r^{7/2}(1-2r)^2$. This concludes the proof.

7.3 Disentangling the tree

Throughout this section, let $x \in V$ and $(x \vee \overline{y} \vee \overline{z})$ be its canonical critical clause. Let L be the left child of the root of T_x and R the right child. Without loss of generality, we

assume that varlabel(L) = y and varlabel(R) = z. Let u, v be two nodes in T_x . If v is a descendant of u and {varlabel(u), varlabel(v)} $\in H_{\text{low}}$, we say v is a subordinate of u and write $v \triangleleft u$.

We fix some $r \in [0, 1]$ and bound $\Pr[\operatorname{Cut}_r(T_x)]$ from below. If $r \ge 1/2$ then $\gamma(r) = 0$ and the events $[\pi(z) < r]$ are independent and have probability r, for all variables. This implies $\Pr[\operatorname{Cut}_r(T_x)] \ge 1 - o_h(1)$ as under the uniform distribution, by Lemma 13. Thus, we focus on r < 1/2.

Removing correlations in T_x . Fix some r < 1/2. To make $\Pr[\operatorname{Cut}_r(T_x)]$ easier to analyze, we take T_x through a series of careful transformation steps, each not increasing $\Pr[\operatorname{Cut}_r(T_x)]$. In every step, we pick a node of v of T_x and give it a fresh label l_v . We have to be careful: if $v \triangleleft u$ for some nodes u in T_x , then variabel(v) and variabel(u) are correlated in D, and this correlation has adverse effects—it decreases $\Pr[\operatorname{Cut}_r(T_x)]$; to take this into account, we have to give the new label a "drag" δ_v . That is, $\Pr[\pi(l_v) < r]$ should not be r but $r - \delta_v$. To state things formally, we define

$$\delta_{\text{root}} := 1.2 \,\epsilon \gamma(r) \max(0, -\phi(r)) \tag{13}$$

$$\delta_{\text{non-root}} := \frac{1.2 \,\epsilon \gamma^2(r)}{1 - r} \tag{14}$$

For every node v of T_x with variabel(v) \notin TwoCC, set

$$\delta_v := \sum_{u:v \triangleleft u} \begin{cases} \delta_{\text{root}} & \text{if } u \text{ is the root} \\ \delta_{\text{non-root}} & \text{if } u \text{ is not the root} \end{cases}$$
(15)

For every v with variabel(v) \in TwoCC, set $\delta_v := -\epsilon \gamma_{\text{TwoCC}}(r)$. We define

$$\delta_{\max} := \max\left(2\delta_{\text{non-root}}, \delta_{\text{non-root}} + \delta_{\text{root}}\right)$$
$$= 1.2 \,\epsilon \gamma(r) \max\left(\frac{2\,\gamma(r)}{1-r}, \frac{\gamma(r)}{1-r} - \phi(r)\right) \,. \tag{16}$$

Since H_{low} has maximum degree 2, the sum in (15) has at most two terms, at most one of them being δ_{root} . It follows that $\delta_v \leq \delta_{\max}$ for all nodes v.

We define DECORR (T_x, r) to be the tree obtained from T_x by assigning a fresh label l_v to each node v in T_x and setting $\Pr[\pi(l_v) < r] := r - \delta_v$. Note that we do not need to fully specify the distribution of $\pi(l_v)$; it is enough to specify $\Pr[\pi(l_v) < r]$ since we are only interested in the event Cut_r , for fixed r. The notation DECORR stands for "de-correlate".

Lemma 46. $\Pr[\operatorname{Cut}_r(T_x)] \ge \Pr[\operatorname{Cut}_r(\operatorname{DeCorr}(T_x, r))].$

The impatient reader may skip the proof and go straight to Section 7.3.1.

Proof. We assign new labels step by step. We start with $T_0 := T_x$.

1. Suppose there is a node v that is a subordinate of two nodes u and w, i.e., $v \triangleleft u$, $v \triangleleft w$. Let a = varlabel(u), b = varlabel(v), c = varlabel(w), so $\{a, b\}, \{b, c\} \in H_{\text{low}}$. We define T_{i+1} by giving v a fresh label b' and sample $\pi(b')$ independently such that $\Pr[\pi(b')] = r - \delta_v$. Looking at the definition of δ_v in (15), we see that

$$\delta_{v} = \begin{cases} 2 \,\delta_{\text{non-root}} & \text{if neither } a \text{ nor } c \text{ is } x \\ \delta_{\text{root}} + \delta_{\text{non-root}} & \text{if } x \in \{a, c\} \end{cases}.$$

Let τ be a random variable that includes the values $\pi(\text{root})$ and, for all labels l except b and b', the information whether $\pi(l) < r$ or $\pi(l) \ge r$.

Claim. $\Pr[\operatorname{Cut}_r(T_i) \mid \tau] \ge \Pr[\operatorname{Cut}_r(T_{i+1}) \mid \tau].$

Proof. Under τ , the event $\operatorname{Cut}_r(T_{i+1})$ becomes a monotone Boolean function $f(z_b, z_{b'})$ in the variables $z_b := \mathbf{1}_{[\pi(b) < r]}$ and $z_{b'} := \mathbf{1}_{[\pi(b') < r]}$. Since T_i can be obtained from T_{i+1} by replacing b' by b, the event $\operatorname{Cut}_r(T_i)$ becomes $f(z_b, z_b)$ under τ . The claim is equivalent to $\Pr[f(z_b, z_b) = 1] \ge \Pr[f(z_b, z_{b'}) = 1]$, all conditioned on τ . We check all possibilities what f could be.

- (a) If $f(z_b, z_{b'})$ does not depend on b' (for example, if $f \equiv 1$, meaning $\operatorname{Cut}_r(T_x)$ happens regardless of $\pi(b), \pi(b')$; or $f \equiv 0$, meaning $\operatorname{Cut}_r(T_x)$ does not happen, regardless of $\pi(b), \pi(b')$; or $f \equiv z_b$, meaning $\operatorname{Cut}_r(T_x)$ happens if and only if $\pi(b) < r$, regardless of $\pi(b')$), then these two probabilities are obviously equal.
- (b) If $f(z_b, z_{b'}) = z_b \wedge z_{b'}$, then obviously

$$\Pr[\operatorname{Cut}_r(T_i) \mid \tau] = \Pr[z_b] \ge \Pr[z_b \land z_{b'}] = \Pr[\operatorname{Cut}_r(T_{i+1}) \mid \tau] .$$

- (c) If $f(z_b, z_{b'}) = z_b \vee z_{b'}$ then... Wait, this cannot happen: by Observation 7, the nodes labeled b and b' form an antichain in T_{i+1} and thus $\operatorname{Cut}_r(T_{i+1})$ cannot become an OR of them.
- (d) This leaves the (most interesting) case that $f(z_b, z_{b'}) = z_{b'}$. We have to show that $\Pr[\pi(b) < r \mid \tau] \ge \Pr[\pi(b') < r \mid \tau]$. The latter is $r \delta_v$ by construction, as $\pi(b')$ is independent of everything.

It remains to show that $\Pr[\pi(b) < r \mid \tau] \ge r - \delta_v$. In the spirit of Corollary 42, we define an interval A_l for each label as: $A_x := \{r\}$; $A_b := [0, r]$; for $l \ne x, b$, set $A_l = [0, r]$ if $\pi(l) < r$ and $A_l = [r, 1]$ if $\pi(l) \ge r$; note that τ contains all necessary information on π to define the A_l . We apply Corollary 42 and conclude that

$$\Pr[\operatorname{Cut}_{r}(T_{i}) \mid \tau] = \Pr[\pi(b) < r \mid \tau]$$

=
$$\Pr[\pi(b) \in A_{b} \mid \pi(l) \in A_{l} \text{ for all } l \neq b]$$

$$\geq r + 1.2 \epsilon \gamma(r) (T_{a}^{-} + T_{c}^{-}) . \qquad (17)$$

Let C be the connected component of H_{low} that contains b. We can apply Corollary 42 since D is independent on all connected components of H_{low} ; that is, conditioning on $\pi(l) \in A_l$ for some $l \notin C$ has no effect and can be ignored. Also, the connected component C of H_{low} is a path or a cycle, so b does not have any neighbors in C besides a and c. Recall that $T_a^- = \mathbb{E}_{s \in A_a}[\phi(s)]$. If $A_a = [0, r]$ then $\mathbb{E}_{s \in A_a}[\phi(s)] = \gamma(r)/r > 0$ and $T_a^- = 0$; If $A_a = [r, 1]$ then $1.2 \epsilon \gamma(r) T_a^- = \frac{-\gamma^2(r)}{1-r} = -\delta_{\text{non-root}}$; if $A_a = \{r\}$ (happens if a = x, the root label of T_x) then $T_a = \phi(r)$ and $1.2 \epsilon \gamma(r) T_a^- = 1.2 \epsilon \gamma(r) \min(\phi(r), 0) = -\delta_{\text{root}}$. This holds analogously for T_c^- . Thus, (17) is at least $r - \delta_v$.

This proves the claim.

We repeat Step 1 until all nodes v with $v \triangleleft u, w$ have received a fresh label.

2. Suppose v is a subordinate of exactly one node $u: v \triangleleft u$. Among all such nodes that have not yet received a fresh label, let v be a minimal such node, i.e., as far from the root as possible. Let a = varlabel(u) and b = varlabel(v). Let c be the other neighbor of b in H_{low} (if there is one; if there isn't, the following proof will only become simpler). Note that no ancestor of v has label c, otherwise $v \triangleleft u, w$ and we would be in Point 1. No descendant of v has label c; otherwise, some wwould have $w \triangleleft v$ and w would fall under Point 1 or under Point 2; the latter case would contradict minimality of v. As before, we assign v a fresh label b' and set $\Pr[\pi(b') < r] = r - \delta_v$, where

$$\delta_v = \begin{cases} \delta_{\text{root}} & \text{if } a = x\\ \delta_{\text{non-root}} & \text{if } a \neq x \end{cases}$$

Let τ be a random variable that includes the values $\pi(\text{root})$ and, for all labels l except b and b' and c, the information whether $\pi(l) < r$ or $\pi(l) \ge r$.

Claim. $\Pr[\operatorname{Cut}_r(T_i) \mid \tau] \ge \Pr[\operatorname{Cut}_r(T_{i+1}) \mid \tau].$

Proof. Under τ , the event $\operatorname{Cut}_r(T_{i+1})|_{\tau}$ becomes a monotone Boolean function $f(z_b, z_{b'}, z_c)$ and $\operatorname{Cut}_r(T_i)|_{\tau}$ becomes $f(z_b, z_b, z_c)$, where $z_l = \mathbf{1}_{[\pi(l) < r]}$. We list all cases what f could be; we skip the trivial cases where $f(z_b, z_{b'}, z_c)$ does not depend on $z_{b'}$.

(a) $f(z_b, z_{b'}, z_c) = z_{b'}$. We define the intervals A_l as in Point 1d; however, for l = c we set $A_c = [0, 1]$ since our condition τ does not reveal anything on c, i.e., we do not condition on c. The rest of the argument is the same, and as in (17) we see that $\Pr[\operatorname{Cut}_r(T_i) \mid \tau]$ is at least

$$\Pr[\operatorname{Cut}_{r}(T_{i}) \mid \tau] = \Pr[\pi(b) < r \mid \tau]$$

$$\geq r + 1.2 \,\epsilon \gamma(r) (T_{a}^{-} + T_{c}^{-})$$

$$= r + 1.2 \,\epsilon \gamma(r) T_{a}^{-} \qquad (\text{since } A_{c} = [0, 1])$$

$$= r - \delta_{v}$$

Thus, $\Pr[\operatorname{Cut}_r(T_i) \mid \tau] \ge r - \delta_v = \Pr[\pi(b') < r] = \Pr[\operatorname{Cut}_r(T_{i+1})].$

- (b) f depends only on b' and c; since no ancestor of u has label c, this means $f(z_{b'}, z_c) = z_{b'} \wedge z_c$. In this case we extend τ to τ' by additionally revealing whether $\pi(c) < r$. If no then then $f(z_{b'}, z_c) = f(z_b, z_c) = 0$, and $\operatorname{Cut}_r(T_i)$, $\operatorname{Cut}_r(T_{i+1})$ both do not happen; if $\pi(c) < r$ then $\Pr[f(z_{b'}, z_c) = 1 | \tau'] = \Pr[\pi(b') < r | \tau']$ and we are back in Point 2a, just with $A_c := [0, r]$; we repeat the above computation, arriving at (17), this time noting that $T_c = \mathbb{E}_{s \in [0,r]}[\phi(s)] = \frac{\gamma(r)}{r} > 0$ and thus $T_c^- = 0$, as well.
- (c) f depends only on $z_{b'}$ and z_b ; then $f(z_b, z_{b'}) = z_b \wedge z_{b'}$ and $\Pr[\operatorname{Cut}_r(T_{i+1}) \mid \tau] = \Pr[z_{b'} \wedge z_b] \leq \Pr[z_b] = \Pr[\operatorname{Cut}_r(T_i)]$ holds trivially.
- (d) f depends on all three. We claim that $f(z_b, z_{b'}, z_c)$ is either $z_{b'} \wedge z_b \wedge z_c$ or $z_{b'} \wedge (z_b \vee z_c)$. More generally, observe that no restriction of f is of the form $z_b \vee z_{b'}$ or $z_{b'} \vee z_c$, since no ancestor or descendant of v has label b or c.

Proposition 47. Let $f(z_1, \ldots, z_n)$ be a monotone Boolean function that depends on z_1 and such that no restriction of f is of the form $z_1 \vee z_i$, for any $i \geq 2$. Then $f(z_1, \ldots, z_n) = z_1 \wedge g(z_2, \ldots, z_n)$ for some monotone Boolean function g.

Proof. Since f is monotone, there is a unique way to write it as a minimal monotone CNF formula $F = C_1 \wedge \ldots C_m$, i.e., such that no clause C of F is contained in another clause D. Since f depends on z_1 , there is a clause containing z_1 , without loss of generality C_1 . If $C_1 = (z_1)$ then no other clause C' of F contains z_1 (by minimality) and $F = z_1 \wedge (C_2 \wedge \cdots \wedge C_m)$, and we are done. Otherwise, without loss of generality $C_1 = (z_1 \vee z_2 \vee \cdots \vee z_k)$. Let ρ be the restriction that sets z_3, \ldots, z_k to 0 and z_{k+1}, \ldots, z_n to 1. Under ρ , the first clause C_1 becomes $z_1 \vee z_2$. Let C_i be any other clause. By minimality, there is some $z_j \in C_i \setminus C_1$, i.e., $j \geq k + 1$. Thus $\rho(z_j) = 1$ and C_j is satisfied. This means that $F|_{\rho} = z_1 \vee z_2$, contradicting the assumption that no restriction of f is $z_1 \vee z_i$.

The proposition shows that $f(z_b, z_{b'}, z_c)$ is either $f(z_b, z_{b'}, z_c)$ is either $z_{b'} \wedge z_b \wedge z_c$ or $z_{b'} \wedge (z_b \vee z_c)$, since we assume that it depends on all three variables. If $f(z_b, z_{b'}, z_c)$ is $z_{b'} \wedge z_b \wedge z_c$ then $\operatorname{Cut}_r(T_i)|_{\tau}$ is $z_b \wedge z_c$ and the claim obviously holds. If $f(z_b, z_{b'}, z_c)$ is $z_{b'} \wedge (z_b \vee z_c)$ then $\operatorname{Cut}_r(T_i)|_{\tau} = f(z_b, z_b, z_c) = z_b \wedge (z_b \vee z_c)$ $z_c) = z_b$. We see that

$$\Pr[\operatorname{Cut}_r(T_{i+1}) \mid \tau] = \Pr[z_{b'} \land (z_b \lor z_c) \mid \tau]$$

$$\leq \Pr[z_{b'} \mid \tau] = r - \delta_v .$$

On the other hand, $\Pr[\operatorname{Cut}_r(T_i)|_{\tau}] = \Pr[z_b \mid \tau] = \Pr[\pi(b) < r \mid \tau]$; we have already seen in Point 2a that this is at least $r - \delta_v$.

This proves the claim.

We repeat Step 2 until every subordinate v has received a new label.

3. Pick any label $b := \text{varlabel}(v) \in V(H_{\text{low}})$ that still appears in T_i . Let a and c be the neighbors of b in H_{low} (again, if b has only one neighbor then pretend c does not exist; the subsequent proof will only become simpler). Neither a nor c appears as the label of an ancestor or descendant of some v with varlabel(b); however, a might appear as an ancestor of c (since a, c are not neighbors in H_{low} , unless a, b, c form a triangle in H_{low}). This makes this step less obvious than one may assume.

We form T_{i+1} by replacing every occurrence of b by a fresh label b' and setting $\Pr[\pi(b') < r] = r$; that is, we set $\delta_v = 0$ for all v with $\operatorname{varlabel}(v) = b$. We let τ contain the values $\pi(l)$ for all labels except a, b, c and b' (yes, now we can afford to reveal the precise value of $\pi(l)$). Under τ , $\operatorname{Cut}_r(T_{i+1})$ becomes a monotone Boolean function $f(z_a, z_{b'}, z_c)$,⁶ and $\operatorname{Cut}_r(T_i)$ becomes $f(z_a, z_b, z_c)$. If $f(z_a, z_b, z_c)$ does not depend on z_b , we are done. Otherwise, since a, c never appear at ancestors or descendants at a b-labeled node, we can apply Proposition 47 and conclude that $f(z_a, z_b, z_c) = z_b \wedge g(z_a, z_c)$.

⁶we replaced all occurrences of b by b', so b does not occur anymore in T_{i+1}

Case 1. $g(z_a, z_c)$ is not $z_a \vee z_c$. Then it is is either 1 or z_a or z_c or $z_a \wedge z_c$ (it cannot be 0 since then f would not depend on z_b , and we have already handled that case). Each of these cases can be equivalently described as $\pi(a) \in A_a$, $\pi(b) \in A_b$, and $\pi(c) \in A_c$ for intervals $A_a A_b, A_c$; for example if $g(z_a, z_c) = z_a \wedge z_c$ then $f(z_a, z_b, z_c) = 1$ if and only if $\pi(a) \in [0, r], \pi(b) \in [0, r]$ and $\pi(c) \in [0, r]$; if $g(z_a, z_c) = z_a$ then $f(z_a, z_b, z_c) = 1$ if and only if $\pi(a) \in [0, r], \pi(b) \in [0, r], \pi(b) \in [0, r]$ and $\pi(c) \in [0, 1]$. In all cases, $A_b = [0, r]$. The condition τ can be described by $\pi(l) \in A_l$ for $A_l := \{\tau(l)\}$. We apply Theorem 40 with $K = \{a, b, c\}$ and I being the remaining variables in their connected component of H_{low} , and see that

$$\Pr[\operatorname{Cut}_{r}(T_{i}) \mid \tau] = \Pr[\pi(k) \in A_{k} \; \forall k \in K \mid \pi(l) \in A_{l} \; \forall i \in I]$$
$$= \mu(A_{a})\mu(A_{b})\mu(A_{c}) \cdot \left(1 + \epsilon \frac{\sum_{\{u,v\} \in E \setminus E_{I}} T_{u}T_{v}}{1 + \epsilon \sum_{\{u,v\} \in E_{I}} T_{u}T_{v}}\right)$$

We partition the set $E \setminus E_I$ into two parts: the first part contains $\{a, b\}$ and $\{b, c\}$, the second part E_{rest} contains the rest. Since a, c are the only neighbors of b in H_{low} , no edge in E_{rest} is incident to b. Furthermore, $\mu(A_b) = \mu([0, r]) = r$, and the terms $T_a T_b$ and $T_b T_c$ are non-negative: each A_u is either [0, r] or [0, 1], and thus T_u is either $\frac{\gamma(r)}{r}$ or 0. Thus, the above is at least

$$r \mu(A_a)\mu(A_c) \cdot \left(1 + \epsilon \frac{T_a T_b + T_b T_c + \sum_{\{u,v\} \in E_{\text{rest}}} T_u T_v}{1 + \epsilon \sum_{\{u,v\} \in E_I} T_u T_v}\right)$$

$$\geq r \mu(A_a)\mu(A_c) \cdot \left(1 + \epsilon \frac{\sum_{\{u,v\} \in E_{\text{rest}}} T_u T_v}{1 + \epsilon \sum_{\{u,v\} \in E_I} T_u T_v}\right)$$
(18)

Next, we compute $\Pr[\operatorname{Cut}_r(T_{i+1}) \mid \tau]$ by observing that $[\pi(b') < r]$ is independent of everything; we set $A_b = [0, 1]$ (since we do not condition on anything; *b* has disappeared), which makes T_b vanish. We apply Theorem 40 and get

$$\Pr[\operatorname{Cut}_{r}(T_{i+1}) \mid \tau] = \Pr[\pi(b') < r] \cdot \Pr[\pi(a) \in A_{a}, \pi(c) \in A_{c} \mid \pi(l) \in A_{l} \forall i \in I]$$
$$= r \cdot \mu(A_{a})\mu(A_{c}) \cdot \left(1 + \epsilon \frac{\sum_{\{u,v\} \in E_{\mathrm{rest}}} T_{u}T_{v}}{1 + \epsilon \sum_{\{u,v\} \in E_{I}} T_{u}T_{v}}\right) = (18) .$$

Thus, $\Pr[\operatorname{Cut}_r(T_i) | \tau] \ge \Pr[\operatorname{Cut}_r(T_{i+1}) | \tau].$

Case 2. $g(z_a, z_c) = z_a \lor z_c$. This is more uncomfortable since $z_b \land (z_a \lor z_c) = 1$ cannot be described as $\pi(a) \in A_a, \pi(b) \in A_b, \pi(c) \in A_c$. At least we know that some node uwith varlabel(u) = a is an ancestor or descendant of some w with varlabel(w) = c. This means that $\{a, c\} \notin H_{\text{low}}$ since we would be in Step 2 otherwise. In other words, a, b, c do not form a triangle in H_{low} , and E_{rest} does not contain $\{a, c\}$. We write

$$z_b \wedge (z_a \vee z_c) = 1 \iff \pi(b) \in [0, r] \wedge \pi(a) \in [0, r] \wedge \pi(c) \in [0, r] \text{ or}$$
$$\pi(b) \in [0, r] \wedge \pi(a) \in [0, r] \wedge \pi(c) \in [r, 1] \text{ or}$$
$$\pi(b) \in [0, r] \wedge \pi(a) \in [r, 1] \wedge \pi(c) \in [0, r]$$

and apply Theorem 40 to each of these three cases. Noting that $T_a = \frac{\gamma(r)}{r}$ if

 $A_a = [0, r]$ and $T_a = \frac{-\gamma(r)}{1-r}$ if $A_a = [r, 1]$, we get

$$\Pr[\operatorname{Cut}_{r}(T_{i}) \mid \tau] = \Pr[\pi(b) < r \land (\pi(a) < r \lor \pi(c) < r) \mid \tau]$$

$$= r^{3} \cdot \left(1 + \epsilon \frac{\frac{\gamma^{2}(r)}{r^{2}} + \frac{\gamma^{2}(r)}{r^{2}} + \sum_{\{u,v\} \in E_{\mathrm{rest}}} T_{u}T_{v}}{1 + \epsilon \sum_{\{u,v\} \in E_{I}} T_{u}T_{v}} \right)$$

$$+ r^{2}(1-r) \cdot \left(1 + \epsilon \frac{\frac{\gamma^{2}(r)}{r^{2}} + \frac{-\gamma^{2}(r)}{r(1-r)} + \sum_{\{u,v\} \in E_{\mathrm{rest}}} T_{u}T_{v}}{1 + \epsilon \sum_{\{u,v\} \in E_{I}} T_{u}T_{v}} \right)$$

$$+ (1-r)r^{2} \cdot \left(1 + \epsilon \frac{\frac{-\gamma^{2}(r)}{r(1-r)} + \frac{\gamma^{2}(r)}{r^{2}} + \sum_{\{u,v\} \in E_{\mathrm{rest}}} T_{u}T_{v}}{1 + \epsilon \sum_{\{u,v\} \in E_{I}} T_{u}T_{v}} \right)$$

Next, we compute $\Pr[\operatorname{Cut}_r(T_{i+1}) | \tau]$, noticing that $\Pr[\pi(b') < r] = r$, independently of everything else, and setting $A_b = [0, 1]$, making T_b vanish; by Theorem 40 we get

$$\Pr[\operatorname{Cut}_{r}(T_{i+1}) \mid \tau] = r \cdot \Pr[\pi(b) \in [0, 1] \land (\pi(a) < r \lor \pi(c) < r) \mid \tau]$$
$$= r^{3} \cdot \left(1 + \epsilon \frac{\sum_{\{u,v\} \in E_{\mathrm{rest}}} T_{u} T_{v}}{1 + \epsilon \sum_{\{u,v\} \in E_{I}} T_{u} T_{v}} \right)$$
$$+ r^{2}(1-r) \cdot \left(1 + \epsilon \frac{\sum_{\{u,v\} \in E_{\mathrm{rest}}} T_{u} T_{v}}{1 + \epsilon \sum_{\{u,v\} \in E_{I}} T_{u} T_{v}} \right)$$
$$+ (1-r)r^{2} \cdot \left(1 + \epsilon \frac{\sum_{\{u,v\} \in E_{\mathrm{rest}}} T_{u} T_{v}}{1 + \epsilon \sum_{\{u,v\} \in E_{I}} T_{u} T_{v}} \right)$$

Taking the difference $\Pr[\operatorname{Cut}_r(T_i) | \tau] - \Pr[\operatorname{Cut}_r(T_{i+1}) | \tau]$, we notice that the 1-term in all three parentheses cancel out, and similarly $\sum_{\{u,v\}\in E_{\text{rest}}} T_u T_v$; the denominator is the same in all cases, and therefore

$$(1 + \epsilon \sum_{\{u,v\} \in E_I} T_u T_v) \cdot (\Pr[\operatorname{Cut}_r(T_i) \mid \tau] - \Pr[\operatorname{Cut}_r(T_{i+1}) \mid \tau])$$

= $r^3 \left(\frac{\gamma^2(r)}{r^2} + \frac{\gamma^2(r)}{r^2} \right) + 2r^2(1-r) \left(\frac{\gamma^2(r)}{r^2} + \frac{-\gamma^2(r)}{r(1-r)} \right)$
= $2(1-r)\gamma^2(r) \ge 0$.

This shows that $\Pr[\operatorname{Cut}_r(T_i) \mid \tau] \geq \Pr[\operatorname{Cut}_r(T_{i+1}) \mid \tau]$. We repeat this step until every $b \in V(H_{\text{low}})$ has been replaced by a new label b'

4. Note that by now no $b \in V(H_{\text{low}})$ appears in T_i . This means that the values $\pi(l)$ are independent for all labels appearing in T_i . Still, some labels might appear multiple times. Let b be a label in T_i that appears multiple times, and pick a node v with varlabel(v) = b. We form T_{i+1} by giving v a new label b' and setting $\Pr[\pi(b') < r] := \Pr[\pi(b) < r]$. Formally, we set $\Pr[\pi(b') < r] = r - \delta_v$ for $\delta_v := r - \Pr[\pi(b) < r]$. Note that δ_v is 0 if $b \notin \text{TwoCC}$, and $\delta_v = -\epsilon \gamma_{\text{TwoCC}}(r)$ if $b \in \text{TwoCC}$.

Let τ contain the values $\pi(l)$ for all labels l except b' and b. We claim that $\Pr[\operatorname{Cut}_r(T_i) \mid \tau] \geq \Pr[\operatorname{Cut}_r(T_{i+1}) \mid \tau]$. Under $\operatorname{Cut}_r(T_{i+1})$ becomes $f(z_b, z_{b'})$ and $\operatorname{Cut}_r(T_i)$ becomes $f(z_b, z_b)$. If $f(z_b, z_{b'})$ is 0 or 1 or z_b then $f(z_b, z_{b'}) = f(z_b, z_b)$. If

 $f(z_b, z_{b'})$ is $z_{b'}$ then

$$\begin{aligned} \Pr[\operatorname{Cut}_r(T_i) \mid \tau] &= \Pr[\pi(b) < r \mid \tau] \\ &= \Pr[\pi(b) < r] & (\text{since } \pi(b) \text{ is independent}) \\ &= \Pr[\pi(b') < r] & (\text{by our choice}) \\ &= \Pr[\pi(b') < r \mid \tau] & (\text{since } \pi(b') \text{ is independent}) \\ &= \Pr[\operatorname{Cut}_r(T_{i+1}) \mid \tau] . \end{aligned}$$

If $f(z_b, z_{b'}) = z_b \wedge z_{b'}$ then $f(z_b, z_{b'}) \leq f(z_b, z_b)$ for all values $z_b, z_{b'}$ and the claim is obvious. Finally, as in all cases above, $f(z_b, z_{b'})$ cannot be $z_b \vee z_{b'}$ since the vertices labeled b in T_i form an antichain.

We repeat this step until no label appears more than once in T_i .

5. Once all nodes except the root have received a new label, stop the process.

Note that the final tree is exactly $DECORR(T_x, r)$ as defined above. This completes the proof of Lemma 46.

Note that Point 4, the case that π is independent of all labels but labels appear multiple times, is exactly the setting Lemma 7 in Paturi, Pudlák, Saks, and Zane [8]. They used the famous FKG inequality; in fact, from Point 4 one can easily extract a proof of the FKG inequality, at least for the special case of monotone Boolean functions.

Of course, we can interrupt the above process anytime, for example if we do want some nodes to keep their shared label, or two variables to keep their correlated distribution. Just that the resulting tree will not be $DECORR(T_x, r)$ but some "intermediate" version thereof.

7.3.1 Making T infinite

For technical reasons, we need some additional clean-up steps. They are mostly dealing with the fact that $DECORR(T_x, r)$ is finite, whereas we would like to deal with infinite trees. We pick some h' with $1 \ll h' \ll h$.

- 1. If some node v with $h' \leq d(v) \leq h-1$ has fewer than two children, add new children until it has two. Make them *safe leaves* if their depth is h and unsafe leaves otherwise. This does not increase $\Pr[\operatorname{Cut}_r]$. Repeat this step for as long as possible.
- 2. Set $\delta_v := \delta_{\max}$ for all nodes v with depth $h' < d(v) \le h$, i.e., $\Pr[\pi(\operatorname{varlabel}(v)) < r] = r \delta_{\max}$. This does not increase $\Pr[\operatorname{Cut}_r]$.

Next, we take the resulting tree and "extend it to infinity". This will increase $\Pr[\operatorname{Cut}_r(T)]$, but only by o(1):

Lemma 48 (Extending T to infinity). Let T be a labeled tree of height h in which all safe leaves have depth h, $\tilde{r} \leq r < 1/2$ be some fixed numbers, and π be a distribution on placements on labels such that for every node v of depth greater than h', the following holds: (1) the label l_v is not shared by any other node; (2) $\pi(l_v)$ is independent of everything else; (3) $\Pr[\pi(l_v) < r] = \tilde{r}$; (4) v has exactly two children if $d(v) \leq h - 1$, and it is a safe leaf if d(v) = h.

Construct T' from T by replacing each safe leaf v by a copy of T_{∞} , so T' has no safe leaves anymore. We set $\Pr[\pi(l) < r] := \tilde{r}$ for all new labels. Then

$$\Pr[\operatorname{Cut}_r(T)] \ge \Pr[\operatorname{Cut}_r(T')] - 2^h r^{h-h'} .$$
(19)

See Lemma C.1 for a proof. One thing to observe is that the lemma does not assume anything about the first h' levels of the tree, neither on the distribution of π on the labels on those layers. Note that the error term in (19) is $(2r)^{h-h'} \cdot 2^{h'}$. Thus, for every fixed r < 1/2, we can find some slowly growing function h' such that $\Pr[\operatorname{Cut}_r(T_x)] \ge$ $\Pr[\operatorname{Cut}_r(T')] - o(1)$. The tree T' is of the following form:

- 1. Up to height h', it looks exactly like DECORR (T_x, r) , and also the distribution of π looks the same on those label.
- 2. A node v of height at least h' has exactly two children.
- 3. $\Pr[\pi(v) < r] \ge r \delta_{\max}$ for all v.
- 4. It has no safe leaves.



From $DeCORR(T_x, r)$ to the partially infinite tree T'

7.3.2 Dealing with bad biases: the biased nodes lemma

We say a label l has a bad bias if $\Pr[\pi(l) < r] = r - \delta_l$ for some $\delta_l > 0$. We need to study how a bad bias at one or several nodes of T decreases the cut probability. For this, let T be an infinite complete binary labeled tree in which all labels are distinct, and let $r \in [0, 1/2]$ be fixed. Let W be a finite set of nodes not containing the root, and suppose we have some $\delta_v \in [0, \delta_{\max}]$ for each $v \in W$, and $\delta_v = 0$ for all $v \in V(T) \setminus W$. Let D be s distribution under which $\Pr[\pi(v) < r] = r - \delta_v$. For a non-root node v, define D_v to be the same distribution, except that $\Pr[\pi(v) < r] = r$. That is, we replace by bias δ_v by 0. It should be clear that $\Pr_D[\operatorname{Cut}_r(T)]$ is smaller than $\Pr_{D_v}[\operatorname{Cut}_r(T)]$. The next lemma shows that it cannot be much smaller:

Lemma 49 (Biased Node lemma). If v is a maximal node in W, i.e., no proper ancestor of v is in W, then

$$\Pr_{D}[\operatorname{Cut}_{r}(T)] \geq \Pr_{D_{v}}[\operatorname{Cut}_{r}(T)] - \delta_{v} \cdot \frac{1 - Q_{r-\delta_{\max}}}{1 - r} \cdot r^{d}$$

where $d := d_T(v)$.

See Lemma C.2 for a proof. Applying the lemma over and over again to all nodes $v \in W$ clears all such nodes of their biases, leaving us with the uniform distribution on all non-root nodes.

Corollary 50 (Biased nodes corollary). If W is finite then

$$\Pr_{D}[\operatorname{Cut}_{r}(T)] \geq Q_{r} - \frac{1 - Q_{r-\delta_{\max}}}{1 - r} \cdot \sum_{v \in V(T) \setminus \{\operatorname{root}\}} \delta_{v} r^{d_{T}(v)}$$

Furthermore, for our particular choices of $\gamma(r)$, $\epsilon \leq 0.13$, and δ_{\max} as defined in (16), it holds that $1 - Q_{r-\delta_{\max}} \leq 1.02 (1 - Q_r)$.

We obtain the "furthermore" part by noting that $Q_{r-\delta_{\max}}$ is decreasing in ϵ and verifying it numerically for $\epsilon = 0.13$.

7.3.3 When $varlabel(u) \in TwoCC$ or u has an only child

Recall the tree T' we have so far. In T', all labels are distinct and $\pi(l)$ is independent for all labels l; T' does not have safe leaves (it is possibly an infinite tree). Every node u has a bias δ_u , meaning $\Pr[\pi(\text{varlabel}(u)) < r] = r - \delta_u$. For a Let u be a node in T'and T_u the subtree of T' rooted at u. Denote by $\operatorname{CLEANSUBTREE}(T', u)$ the tree obtained from T' by replacing T_u with a copy of T^{∞} , and setting $\delta_w = 0$ for all nodes in that new copy. In this section, we will show that $\Pr[\operatorname{Cut}_r(T')] \ge \Pr[\operatorname{Cut}_r(\operatorname{CLEANSUBTREE}(u, T'))]$ if (1) u has at most one child or (2) $\delta_u = -\epsilon \gamma_{\operatorname{TwoCC}}(r)$. We apply this cleanup step over and over again as long as there is a node satisfying condition (1) or (2); we denote the resulting tree by T''. Recall the definition of $\operatorname{Can}(T_x)$, the set of canonical nodes of T_x ; every non-canonical node v of T_x has an ancestor u such that u has at most one child or varlabel $(u) \in \operatorname{TwoCC}$. This means that all non-canonical nodes v will be replaced by this clean-up step and thus $\delta_v = 0$ in T''.

Proposition 51 (TwoCC cleanup). If T' is as above and $\delta_u = -\epsilon \gamma_{\text{TwoCC}}(r)$, then $\Pr[\text{Cut}_r(T')] \ge \Pr[\text{Cut}_r(\text{CLEANSUBTREE}(u,T'))]$, for our specific choices of $\gamma_{\text{TwoCC}}(r) = 40r^{7/2}(1-2r)^2$, $\gamma(r) = r(1-2r)^{3/2}$, and $\epsilon \le 0.13$.

See Proposition C.3 for a proof. Overall, the exact choice of γ_{TwoCC} is not very crucial. It must be large enough to satisfy (44); besides that, it should not be too large: indeed, γ_{TwoCC} being large means that, under D, variables $a \in \text{TwoCC}$ tend to come early; this depresses $\Pr[\text{Forced}(a)]$; this is okay since TwoCC-variables have a "natural advantage" by virtue of their having two critical clauses; however, if γ_{TwoCC} is too large, this advantage will be all eaten up. Also, γ_{TwoCC} should not be too *steep*, since otherwise $\operatorname{KL}(D_{\epsilon}^{\gamma_{\text{TwoCC}}}||U)$ becomes too large. Our choice of γ_{TwoCC} is simply a hand-crafted polynomial that ticks both boxes, and is pretty arbitrary beyond that. Next, we deal with nodes that have only one child—or none at all, i.e., unsafe leaves.

Proposition 52 (One-child cleanup). Suppose T' is as above and u has at most child. Furthermore, suppose that $\delta_u \leq \delta_{\max}$ for all nodes u, and $r(1-2r) \geq 2\delta_{\max}$. Then $\Pr[\operatorname{Cut}_r(T')] \geq \Pr[\operatorname{Cut}_r(\operatorname{CLEANSUBTREE}(u,T'))].$

See Proposition C.4 for a proof. Note that the condition $r(1-2r) \ge 2\delta_{\max}$ holds for our choice of $\gamma = r(1-2r)^{3/2}$ and $\epsilon \le 0.13$ (for what it's worth, it even holds for all $\epsilon \le 1$). We apply Propositions 51 and 52 wherever applicable (strictly speaking, it is enough to apply it to every *maximal* node *u* among those with varlabel(*u*) \in TwoCC or *u* having an only child). Denote the resulting tree by CLEANUP(T_x, r). Let us pause for a minute and look at the shape of T'. It is an infinite binary tree, and all nodes have distinct labels. The event [π (varlabel(v)) < r] has probability $r - \delta_v$, where δ_v is as defined in (15) if v is a canonical node and $\delta_v = 0$ if v is a non-canonical node or a new node (meaning newly created in the process of CLEANUP(T_x, r)).

7.4 General but pessimistic lower bound on $Pr[Cut(T_x)]$

Let $T'_r := \text{CLEANUP}(T_x, r).$

Lemma 53. If $x \notin \text{TwoCC}$ then $\Pr_D[\text{Cut}(T_x)] \ge \int_0^1 \Pr_D[\text{Cut}_r(T'_r)] \ge s_3 - o(1) - 1.1 \epsilon \text{Thr.}$

Proof. For now, fix some r < 1/2 and write $T' := \text{CLEANUP}(T_x, r)$.⁷ The tree T' and the distribution on its labels is sufficiently simple so we can analyze $\Pr[\text{Cut}_r(T')]$. Note that $\delta_v > 0$ only if v is a canonical node of T_x , where δ_v is as defined in (15). There are only finitely many canonical nodes (since T_x itself is finite), and thus we can apply Corollary 50. Writing $d(v) := d_{T_x}(v)$ to denote the depth of a node v in T_x , we get

$$\Pr_{D}[\operatorname{Cut}_{r}(T')] \geq Q_{r} - 1.02 \frac{1 - Q_{r}}{1 - r} \sum_{\substack{v \in \operatorname{Can}(T_{x}) \\ v \text{ not root}}} \delta_{v} r^{d(v)}$$

$$= Q_{r} - 1.02 \frac{1 - 2r}{(1 - r)^{3}} \sum_{\substack{v \in \operatorname{Can}(T_{x}) \\ v < \text{root}}} \delta_{\operatorname{root}} r^{d(v)}$$

$$- 1.02 \frac{1 - 2r}{(1 - r)^{3}} \sum_{\substack{v \in \operatorname{Can}(T_{x}) \\ v < \text{root}}} \sum_{\substack{u \in V(T_{x}) \\ v < u \text{ not root}}} \delta_{\operatorname{non-root}} r^{d(v)} .$$

$$(20)$$

Let us work on the first sum. Recall that by definition, a node $v \in \operatorname{Can}(T_x)$ has $v \triangleleft \operatorname{root}(T_x)$ if $\{\operatorname{varlabel}(v), x\} \in H_{\text{low}}$. Writing $x' \sim x$ as a shorthand for $\{x', x\} \in H_{\text{low}}$, we get

$$1.02 \frac{1-2r}{(1-r)^3} \sum_{v \in \operatorname{Can}(T_x), v < \operatorname{root}} \delta_{\operatorname{root}} r^{d(v)}$$

$$= 1.02 \frac{1-2r}{(1-r)^3} \sum_{x' \sim x} \sum_{v: \operatorname{varlabel}(v) = x'} \delta_{\operatorname{root}} r^{d(v)}$$

$$= \frac{1.02 \delta_{\operatorname{root}}}{1-2r} \cdot \sum_{x' \sim x} \frac{(1-2r)^2}{(1-r)^3} \sum_{v: \operatorname{varlabel}(v) = x'} r^{d(v)}$$

$$= \frac{1.02 \delta_{\operatorname{root}}}{r(1-2r)} \cdot \sum_{x' \sim x} \frac{(1-2r)^2}{(1-r)^3} \sum_{v: \operatorname{varlabel}(v) = x'} r^{d(v)+1}$$

$$= \frac{1.02 \delta_{\operatorname{root}}}{r(1-2r)} \sum_{x' \sim x} \operatorname{LABELDENSITY}(x', T_x, r) \qquad \text{(as defined in (9))}$$

$$\leq 0.28 \epsilon \sum_{x' \sim x} \operatorname{LABELDENSITY}(x', T_x, r)$$

where we numerically verify that $\frac{1.02 \,\delta_{\text{root}}}{r(1-2 \,r)} \leq 0.28 \,\epsilon$ for all r. The sum in (20) is slightly more subtle to bound:

$$1.02 \frac{1-2r}{(1-r)^3} \sum_{\substack{v \in \operatorname{Can}(T_x) \\ v \text{ not root } u \text{ not root} \\ v < v < v < u}} \sum_{\substack{u \in V(T_x) \\ v < v < u}} \delta_{\operatorname{non-root}} r^{d(v)}$$
$$= 1.02 \frac{1-2r}{(1-r)^3} \sum_{\substack{u \in \operatorname{Can}(T_x) \\ u \text{ not root }}} \sum_{\substack{v \in V(T_x) \\ v < u}} \delta_{\operatorname{non-root}} r^{d(v)}$$
$$= \frac{1.02 \delta_{\operatorname{non-root}}}{1-2r} \sum_{\substack{u \in \operatorname{Can}(T_x) \\ u \text{ not root }}} r^{d(u)} \cdot \frac{(1-2r)^2}{(1-r)^3} \sum_{\substack{v \in \operatorname{Can}(T_x), v < u}} r^{d(v)-d(u)}$$

⁷For $r \ge 1/2$ we already know that $\Pr[\operatorname{Cut}_r(T')] = 1$.
We apply Lemma 33 to observe that summing over canonical nodes $v \triangleleft u$ is at most as much as summing over $v' \in T_a$, where T_a is the canonical clause tree of a := varlabel(u). The above sum is at most

$$\frac{1.02 \,\delta_{\text{non-root}}}{1-2 \,r} \sum_{\substack{u \in \text{Can}(T_x)\\u \text{ not root}\\a:=\text{varlabel}(u)}} r^{d(u)} \cdot \frac{(1-2 \,r)^2}{(1-r)^3} \sum_{\substack{v' \in \text{Can}(T_a)\\v' < \text{root}(T_a)}} r^{d_{T_a}(v)}$$

$$= \frac{1.02 \,\delta_{\text{non-root}}}{1-2 \,r} \sum_{\substack{u \in \text{Can}(T_x)\\u \text{ not root}\\a:=\text{varlabel}(u)}} r^{d(u)-1} \cdot \sum_{b\sim a} \frac{(1-2 \,r)^2}{(1-r)^3} \sum_{\substack{v' \in \text{Can}(T_a)\\varlabel(v')=b}} r^{d_{T_a}(v)+1}$$

$$= \frac{1.02 \,\delta_{\text{non-root}}}{1-2 \,r} \sum_{\substack{u \in \text{Can}(T_x)\\u \text{ not root}\\a:=\text{varlabel}(u)}} r^{d(u)-1} \cdot \sum_{b\sim a} \text{LABELDENSITY}(b, T_a, r) .$$

where LABELDENSITY (x_i, T_x, r) is as defined in (9).

Proposition 54. $\frac{1.02\,\delta_{\text{non-root}}}{1-2\,r}r^{d-1} \le 4.896\,\epsilon \left(\frac{1}{2}\right)^d \frac{(d+1)^{d+1}}{(d+3)^{d+3}} \text{ for all } r \in [0, 1/2].$

We prove this as Proposition C.5 in the appendix. Putting things together, we conclude that

$$\Pr[\operatorname{Cut}_{r}(T')] \geq Q_{r} - 0.28 \epsilon \sum_{x' \sim x} \operatorname{LABELDENSITY}(x', T_{x}, r) \\ - 4.896 \epsilon \sum_{\substack{u \in \operatorname{Can}(T_{x}) \\ u \text{ not root} \\ a := \operatorname{varlabel}(u)}} \left(\frac{1}{2}\right)^{d(u)} \frac{(d(u) + 1)^{d(u) + 1}}{(d(u) + 3)^{d(u) + 3}} \sum_{b \sim a} \operatorname{LABELDENSITY}(b, T_{a}, r)$$

Note that the marginal distribution of $\pi(x)$ is uniform under D, and $\pi(x)$ is independent of everything else in T'; we get $\Pr[\operatorname{Cut}(T_x)]$ by integrating over r (we will drop the o(1)for notational convenience):

$$\Pr[\operatorname{Cut}(T_x)] \ge \int_0^1 \Pr[\operatorname{Cut}_r(T_x)] \, dr \ge \int_0^1 \Pr[\operatorname{Cut}_r(T')] \, dr - o(1)$$

$$\ge s_3 - 0.28 \, \epsilon \sum_{\substack{x' \sim x}} \text{LABELDENSITY}(x', T_x)$$

$$- 4.896 \, \epsilon \sum_{\substack{u \in \operatorname{Can}(T_x)\\ u \text{ not root}\\ a := \text{varlabel}(u)} \left(\frac{1}{2}\right)^{d(u)} \frac{(d(u) + 1)^{d(u) + 1}}{(d(u) + 3)^{d(u) + 3}} \sum_{b \sim a} \text{LABELDENSITY}(b, T_a)$$

Note that $\{a, b\} \in H_{\text{low}}$ means, by construction, that LABELDENSITY $(b, T_a) \leq \text{THR}$, and similar for $\{x', x\} \in H_{\text{low}}$. Therefore, the terms LABELDENSITY (\cdot, T_{\cdot}) in the above expression are all at most THR. Furthermore, there are at most two variables x' with $x' \sim x$, and at most two b with $b \sim a$ (note that H_{low} has maximum degree 2), meaning the first sum and the inner sum of the second sum are at most 2 THR each. Since for each integer $d \geq 1$ there are at most 2^d nodes $u \in \text{Can}(T_x)$ of depth d, we get

$$\Pr[\operatorname{Cut}(T_x)] \ge s_3 - 0.56 \,\epsilon \operatorname{Thr} - 9.792 \,\epsilon \operatorname{Thr} \sum_{d=1}^{\infty} \frac{(d+1)^{d+1}}{(d+3)^{d+3}}$$
$$\ge s_3 - 0.56 \,\epsilon \operatorname{Thr} - 0.54 \,\epsilon \operatorname{Thr} \qquad \text{(see Proposition C.6)}$$
$$= s_3 - 1.1 \,\epsilon \operatorname{Thr} .$$

This concludes the proof of Lemma 53.

7.5 The case that $x \notin \text{TwoCC}$ and $\{y, z\} \in H_{\text{low}}$

This is arguably the heart of this paper, where the improvement from introducing correlations between variables shows up. With the work done in the previous sections, this will go rather painlessly.

Lemma 55. For our choice $\gamma(r) = r(1-2r)^{3/2}$ and $\epsilon \leq 0.13$, it holds that $\Pr_D[\operatorname{Cut}(T_x)] \geq s_3 - o(1) - 1.1 \epsilon \operatorname{Thr} + \epsilon \int_0^{1/2} \gamma^2(r)(1-Q_r)^2 \geq s_3 - o(1) - 1.1 \epsilon \operatorname{Thr} + 0.001687 \epsilon$.

Proof. As before, let $(x \vee \overline{y} \vee \overline{z})$ be the critical clause of x, and L and R be the left and right children $\operatorname{root}(T_x)$, respectively. Without loss of generality, $\operatorname{varlabel}(L) = y$ and $\operatorname{varlabel}(R) = z$. Let T_L be the subtree of T_x rooted at L and T_R the one rooted at R.

Proposition 56. If $\{y, z\} \in H_{\text{low}}$, then $\{x, y\}, \{x, z\} \notin H_{\text{low}}$.

Proof. Suppose, for the sake of contradiction, that $\{x, y\} \in H_{\text{low}}$. This means that there is a variable a such that $(a \lor \bar{x} \lor \bar{y})$ is its canonical critical clause. If $a \neq z$, then this clause, together with $(x \lor \bar{y} \lor \bar{z})$ implies the clause $(a \lor \bar{y} \lor \bar{z})$, meaning that a has two critical clauses (recall that we included these "derived" critical clauses in our definition of TwoCC). This is a contradiction because it would remove the $\{x, y\}$ from H_{free} , thus it would never end up in H_{low} .

If a = z then F contains $C_x = (x \vee \bar{y} \vee \bar{z})$ and $C_z = (z \vee \bar{x} \vee \bar{y})$. Since $(1, \ldots, 1)$ is the unique satisfying assignment, the assignment $\alpha[x = z = 0]$ must violate some clause distinct from C_x and C_z ; if this is a critical clause for x, then x has two critical clause, a contradiction as observed above; if it is critical for z, then z has two critical clauses, also leading to $\{y, z\}$ being removed from H_{free} . It could be a clause of the form $(x \vee z \vee \bar{b})$. In the latter case, $(x \vee z \vee \bar{b})$ and $(x \vee \bar{y} \vee \bar{z})$ together imply the clause $(x \vee \bar{y} \vee \bar{b})$, whether y = b or not, and again x has two critical clauses.

The upshot is that x is not a neighbor of y nor z in H_{low} , which implies the following: **Proposition 57.** The marginal distribution of $(\pi(y), \pi(z))$, conditioned on $\pi(x) = r$, is still $D_{\epsilon}^{\gamma, \Box}$:

See Proposition C.7 for a proof. For fixed r, write $T'' = T''(r) := \text{CLEANUP}(T_x, r)$, the same as defined in the last section (although it is called T' in the last section). We could of course argue as before and conclude that

$$\Pr_{D}[\operatorname{Cut}(T_{x})] \ge \int_{0}^{1} \Pr_{D}[\operatorname{Cut}_{r}(T'')] \, dr - o(1) \ge s_{3} - o(1) - 1.1 \, \epsilon \mathrm{THR}$$

but this would defeat its purpose since we want to show that $\Pr_D[\operatorname{Cut}(T_x)] > s_3$, and by a significant margin. Instead, we define a "partially cleaned-up tree" T' by taking T'' and putting the variables y, z back as labels of L and R. To put it differently, we run the procedure $\text{CLEANUP}(T_x)$ but, while executing DECORR, we stop short of assigning L and R a fresh label. It takes a minute of thought to see that this is still "legal", i.e., $\Pr_D[\operatorname{Cut}_r(T_x)] \ge \Pr_D[\operatorname{Cut}_r(T')] - o(1).$

In T', the values $\pi(l)$ are independent for all labels except y, z. In particular, the events $\operatorname{Cut}_r(T'_L)$ and $\operatorname{Cut}_r(T'_R)$ are independent (note that $\operatorname{Cut}_r(T)$ does not depend on the label of the root of T, whereas w $\operatorname{Cut}_r(T)$ does). Also, they are independent of the pair $(\pi(y), \pi(z))$. Therefore,

$$\begin{split} \Pr_{D}[\operatorname{Cut}_{r}(T')] &= & \Pr_{D}[\operatorname{Cut}_{r}(T'_{L}) \wedge \operatorname{Cut}_{r}(T'_{R}] \\ &+ & \Pr_{D}[\neg \operatorname{Cut}_{r}(T'_{L}) \wedge \operatorname{Cut}_{r}(T'_{R})] \cdot \Pr[\pi(y) < r] \\ &+ & \Pr_{D}[\operatorname{Cut}_{r}(T'_{L}) \wedge \neg \operatorname{Cut}_{r}(T'_{R})] \cdot \Pr[\pi(z) < r] \\ &+ & \Pr_{D}[\neg \operatorname{Cut}_{r}(T'_{L}) \wedge \neg \operatorname{Cut}_{r}(T'_{R})] \cdot \Pr[\pi(y), \pi(z) < r] \\ &= & \Pr_{D}[\operatorname{Cut}_{r}(T'_{L}) \wedge \operatorname{Cut}_{r}(T'_{R})] \\ &+ & \Pr_{D}[\neg \operatorname{Cut}_{r}(T'_{L}) \wedge \operatorname{Cut}_{r}(T'_{R})] \cdot r \\ &+ & \Pr_{D}[\operatorname{Cut}_{r}(T'_{L}) \wedge \neg \operatorname{Cut}_{r}(T'_{R})] \cdot r \\ &+ & \Pr_{D}[\neg \operatorname{Cut}_{r}(T'_{L}) \wedge \neg \operatorname{Cut}_{r}(T'_{R})] \cdot (r^{2} + \epsilon \gamma^{2}(r) \end{split}$$

by Proposition 39. Since $\pi(l)$ is independent on all labels in $T'' = \text{CLEANUP}(T_x)$, the fully cleaned-up tree, we see that

$$\begin{aligned} \Pr_{D}[\operatorname{Cut}_{r}(T'')] &= & \Pr_{D}[\operatorname{Cut}_{r}(T''_{L}) \wedge \operatorname{Cut}_{r}(T''_{R}] \\ &+ & \Pr_{D}[\neg \operatorname{Cut}_{r}(T''_{L}) \wedge \operatorname{Cut}_{r}(T''_{R})] \cdot r \\ &+ & \Pr_{D}[\operatorname{Cut}_{r}(T''_{L}) \wedge \neg \operatorname{Cut}_{r}(T''_{R})] \cdot r \\ &+ & \Pr_{D}[\neg \operatorname{Cut}_{r}(T''_{L}) \wedge \neg \operatorname{Cut}_{r}(T''_{R})] \cdot r^{2} \end{aligned}$$

Note that $\operatorname{Cut}_r(T'_L)$ and $\operatorname{Cut}_r(T''_L)$ the same events, since T'_L and T''_L differ only in the label of their root (y in the case of T'_L , some fresh label in the case of T''_L), and similarly $\operatorname{Cut}_r(T'_R) = \operatorname{Cut}_r(T'_R)$. Thus,

$$\Pr_{D}[\operatorname{Cut}_{r}(T')] = \Pr_{D}[\operatorname{Cut}_{r}(T'')] + \epsilon \gamma^{2} \Pr_{D}[\neg \operatorname{Cut}_{r}(T''_{L}) \land \neg \operatorname{Cut}_{r}(T''_{R})] .$$

The trees T''_L, T''_R have no leaves (they are infinite binary trees), all their have distinct labels, π is independent on them, and $\Pr[\pi(\text{varlabel}(v)) < r] = r - \delta_v \leq r$ for all nodes vin them. Therefore $\Pr[\text{Cut}_r(T''_L)], \Pr[\text{Cut}_r(T''_L)] \leq Q_r$ and

$$\Pr_{D}[\operatorname{Cut}_{r}(T') \ge \Pr_{D}[\operatorname{Cut}_{r}(T'')] + \epsilon \gamma^{2} (1 - Q_{r})^{2}$$

Integrating over $r = \pi(x)$, we get

$$\Pr_{D}[\operatorname{Cut}(T_{x})] \geq \int_{0}^{1} \Pr_{D}[\operatorname{Cut}_{r}(T'')] dr + \epsilon \int_{0}^{1/2} \gamma^{2} (1 - Q_{r})^{2}$$
$$\geq s_{3} - o(1) - 1.1 \epsilon \operatorname{THR} + \epsilon \int_{0}^{1/2} \gamma^{2} (r) (1 - Q_{r})^{2} ,$$

by Lemma 53 from the previous section (observe that our T'' here is the same as the $T' = \text{CLEANUP}(T_x, r)$ from the last section). This concludes the proof of Lemma 55. \Box

7.6 The case that $x \notin \text{TwoCC}$ and $\{y, z\} \in H_{\text{high}}$

For technical reasons, let T_x be a critical clause tree of height h', where h' is sufficiently large compared to h, but still a slowly growing function in n.

Lemma 58. $\Pr_D[\operatorname{Cut}(T_x)] \ge s_3 - o(1) - 1.1 \epsilon \operatorname{THR} + 0.9 \operatorname{THR}.$

Intuition. Without loss of generality, LABELDENSITY $(z, T_y) \geq$ THR. That is, many nodes of T_y have label z and they are close to the root. We will show that this means, in general, that T_x has not only a child labeled z (the right child), but z occurs also very often in the left subtree of T_x . This gives a significant boost to $\text{Cut}(T_x)$.

Proof. Let L and R be the left and right child of the root of T_x , respectively, such that $\operatorname{varlabel}(L) = y$ and $\operatorname{varlabel}(R) = z$. Let T_L and T_R be the subtrees of T_x rooted at L and R, respectively. If $\{y, z\} \in H_{\operatorname{high}}$, then without loss of generality, z is the *dense* variable of $\{y, z\}$, meaning LABELDENSITY $(z, T_y) \geq \operatorname{THR}$ and $z \notin V(H_{\operatorname{low}})$. Thus, $\pi(z)$ is independent of everything else under D. Let A_z be the set of $u \in \operatorname{Can}(T_y)$ that carry label z. We would like to argue that we can "find a copy" of A_z in T_L . After all, T_L should look a lot like T_y . This is, however, not true in general, but morally, it can only fail in our favor. To be more precise:

Proposition 59 (Copy of A_z in T_L). There is an antichain A'_z of nodes in T_y that is "above A_z ", i.e., for every $u \in A_z$, there is $u' \in A'_z$ such u' is a (not necessarily proper) ancestor of u.

Additionally, there are disjoint sets $B_z, B_1 \subseteq V(T_L)$ such that $B_z \cup B_1$ is an antichain, and a bijection $\Phi : A'_z \to (B_z \cup B_1)$ such that for every $u \in A'_z$:

- 1. the T_y -path from $root(T_y)$ to u and the T_L -path from L to $\Phi(u)$ have the same length and the same label sequence;
- 2. if $\Phi(u) \in B_z$ then variabel $(u) = \text{variabel}(\Phi(u)) = z$;
- 3. if $\Phi(u) \in B_1$ then variabel(u) = variabel $(\Phi(u)) \neq z$ and $\Phi(u)$ has at most one child in T_x .

Furthermore, for all $v \in B_1 \cup B_z$, all proper ancestors of v are canonical.

Proof. For each $u \in A_z \subseteq V(T_y)$, we try to find a corresponding $u' \in V(T_L)$. To do this, walk along the path p from $\operatorname{root}(T_y)$ to u, and let $y = l_0, l_1, \ldots, l_d = \operatorname{varlabel}(u)$ be the labels along this path. Try to find a corresponding path p' in T_x , starting at L, with the same label sequence. Formally, let d' be maximum such that T_x contains a path p' of canonical nodes starting at L with label sequence $l_0, l_1, \ldots, l_{d'}$. Such a d' exists since $\operatorname{varlabel}(L) = y = l_0$.

If d' = d then p' is a "copy" of p in T_L . We let u' be its endpoint, add u to A'_z , u' to B_z , and set $\Phi(u) = u'$.

If d' < d, let v be the $(d')^{\text{th}}$ node on p (start counting at 0) and let v' be the endpoint of p'. By assumption, v is a canonical node in T_y ; let $a := \text{varlabel}(v) = l_{d'}$ and $(a \lor \overline{b} \lor \overline{c})$ be its canonical critical clause. Without loss of generality, $l_{d'+1} = b$. The key question is: why did the construction of T_x not use the clause $(a \lor \overline{b} \lor \overline{c})$ as clause label of v'? The only reason can be that $\alpha_{v'}$, the assignment label of node v', satisfies it; but α_v violates it, by construction of T_y . Since $\alpha_{v'}$ and α_v differ only in $\alpha_{v'}(x) = 0$ and $\alpha_v(x) = 1$. Thus, $x \in \{b, c\}$. What is clauselabel(v')? It cannot be a critical clause, otherwise a would have two critical clauses, putting it in TwoCC, contradicting the assumption that v is a canonical node. Thus, it must be some non-critical clause, meaning that v' has at most one child. We put v into A'_z and v' into B_1 and set $\Phi(v) = v'$.

This figure illustrates the statement of Proposition 59.



From now on, throughout this section, fix the set B_z, B_1 as in the proposition. They will allow us to argue that $\Pr_D[\operatorname{Cut}(T_x)]$ is significantly larger than s_3 . As in the previous section, jumping right away to $T'' := \operatorname{CLEANUP}(T_x, r)$ would delete what we gain from B_z and B_1 . We have to define a "partially cleaned up version" T': we apply the procedure $\operatorname{CLEANUP}(T_x)$ but

- 1. If $v \in B_1$, let w be its only child. Do clean up the subtree T_w but do not add another child to v; v should still have only one child in T'. Do assign v a fresh label l with $\Pr_D[\pi(l) < r] = r - \delta_v$, though.
- 2. if $v \in B_z$, leave its label z; don't assign a fresh label.

Observation 60. Let T' be the result of the partial cleanup procedure. Then $\Pr_D[\operatorname{Cut}(T_x)] \ge \Pr_D[T'] - o(1)$.

Let r be fixed. When analyzing $\Pr[\operatorname{Cut}_r(T')]$, we want to quantify by how much a node with an only child or an additional node with label z (besides R) boosts $\operatorname{Cut}_r(T')$. To this end, we define a "one-child bonus" and a "multiple-label bonus":

$$OCB(d,r) := \frac{r(1-2r) - 2\delta_{\max}(1-r)}{(1-r)^2} \cdot \left(r - \frac{\delta_{\max}}{1-r}\right)^d , \qquad (21)$$

$$MLB(d,r) := \frac{(1-2r)^2 r}{(1-r)^3} \cdot \left(r - \frac{\delta_{\max}}{1-r}\right)^{d-1} .$$
(22)

Lemma 61. Let r < 1/2 and $T'' := \text{CLEANUP}(T_x, r)$. Then

$$\operatorname{Cut}_r(T') \ge \operatorname{Cut}_r(T'') + \sum_{v \in B_1} \operatorname{OCB}(d_{T_x}(v), r) + \sum_{v \in B_z} \operatorname{MLB}(d_{T_x}(v), r) .$$

Proof. Order the elements in B_1 arbitrarily; set $T_0 := T'$; for $1 \le i \le |B_1|$, construct T_i from T_{i-1} as follows: let u be the ith node of B_1 ; replace T_u , the subtree of T_{i-1} rooted at u, by $T_{u'}$, a fresh copy of T^{∞} , and set $\delta_w = 0$ for all nodes in $T_{u'}$, i.e., $\Pr[\pi(w) < r] = r$.

Lemma 62 (Bonus from only child). Let $d = d_{T_x}(u)$ be the distance from the root to u. Then $\Pr[\operatorname{Cut}_r(T_{i-1})] \ge \Pr[\operatorname{Cut}_r(T_i)] + \operatorname{OCB}(d, r)$.

This is proved as Lemma C.8 in the appendix. Applying the lemma for $i = 1, ..., |B_1|$, we conclude that

$$\Pr_{D}[\operatorname{Cut}_{r}(T_{|B_{1}|})] \ge \Pr_{D}[\operatorname{Cut}_{r}(T')] + \sum_{v \in B_{1}} \operatorname{OCB}(d) .$$
(23)

What does T_{B_1} look like? It is an infinite binary tree, all nodes have two children, and all nodes have distinct labels except those in $B_z \cup \{R\}$, which are all labeled z. All non-canonical nodes and all newly introduced nodes have $\delta_v = 0$, i.e., $\Pr[\pi(v) < r] = r$.

Next, we get rid of the multiple labels z: Order the elements of B_z arbitrary and let $T'_0 := T_{B_1}$. For $|B_1| + 1 \le i \le |B_1| + |B_z|$, let u be the $(i - |B_1|)^{\text{th}}$ element of B_z and construct T'_i from T'_{i-1} by giving u a fresh label l, and set $\delta_u = 0$, i.e., $\Pr[\pi(l) < r] = r$. Let $T'' := T_{|B_1| + |B_z|}$ be the final result of this procedure.

Lemma 63 (Bonus from multiple labels). Let $d = d_{T_x}(u)$ be the distance from the root to u. Then $\Pr[\operatorname{Cut}_r(T'_{i-1})] \ge \Pr[\operatorname{Cut}_r(T'_i)] + \operatorname{MLB}(d)$.

See Lemma C.9 for a proof. Applying Lemma 63 for $i = |B_1| + 1, \ldots, |B_1| + |B_2|$ and combining this with (23), we get $T'' := T_{|B_1|+|B_2|}$. This completes the work of CLEANUP(T_x) since steps $i = 1, \ldots, |B_1| + |B_2|$ just described simply are the steps that we skipped in our partial cleanup version T'. In other words, $T'' = \text{CLEANUP}(T_x)$. Altogether,

$$\Pr_D[\operatorname{Cut}_r(T')] \ge \Pr_D[\operatorname{Cut}_r(T'')] + \sum_{v \in B_1} \operatorname{OCB}(d_{T_x}(v), r) + \sum_{v \in B_z} \operatorname{MLB}(d_{T_x}(v), r) .$$

This concludes the proof of Lemma 61.

Integrating over all r, we get

$$\begin{aligned} \Pr_{D}[\operatorname{Cut}(T_{x})] &\geq \int_{0}^{1} \left(\Pr_{D}[\operatorname{Cut}_{r}(T'')] + \sum_{v \in B_{1}} \operatorname{OCB}(d(v), r) + \sum_{v \in B_{z}} \operatorname{MLB}(d(v), r) \right) - o(1) \\ &\geq s_{3} - o(1) - 1.1 \,\epsilon \operatorname{THR} + \sum_{v \in B_{1}} \int_{0}^{1/2} \operatorname{OCB}(d(v), r) \, dr + \sum_{v \in B_{z}} \int_{0}^{1/2} \operatorname{MLB}(d(v), r) \, dr \\ &= s_{3} - o(1) - 1.1 \,\epsilon \operatorname{THR} + \sum_{v \in B_{1}} \operatorname{OCB}(d(v)) + \sum_{v \in B_{z}} \operatorname{MLB}(d(v)) \end{aligned}$$

using Lemma 53. It remains to bound the sums from below.

Lemma 64. $\sum_{v \in B_1} \text{OCB}(d(v)) + \sum_{v \in B_z} \text{MLB}(d(v)) \ge 0.9 \text{ Thr.}$

See Lemma C.10 for a proof. We conclude that $\Pr_D[\operatorname{Cut}(T_x)] \ge s_3 - o(1) - 1.1 \epsilon \operatorname{Thr} + 0.9 \operatorname{Thr}$, which proves Lemma 58.

7.7 The case that $x \in \text{TwoCC}$

Hansen, Kaplan, Zamir, and Zwick [3] proved that if $x \in \text{TwoCC}$, then we can build a critical clause tree T_x for which $\Pr_U[\text{Cut}_r(T_x)] \ge Q_r \cdot B(r) - o_h(1)$ for every $r \in [0, 1]$, where

$$B(r) := \begin{cases} 1 + \frac{(1-2r)^2(1-2r+2r^2)}{(1-r)^2} & \text{if } r < 1/2.\\ 1 & \text{if } r \ge 1/2 \end{cases}$$
(24)

This holds under the uniform distribution. Under D, let us go through the labels of T_x in some order, starting with x, namely $x = l_0, l_1, \ldots$ Let

$$\tau_i := (\pi(x), \mathbf{1}_{[\pi(l_1) < r]}, \dots, \mathbf{1}_{[\pi(l_i) < 1]})$$

For a concrete choice of τ_i , what is $\Pr[\pi(l_{i+1}) < r \mid \tau_i]$? If $l_{i+1} \notin V(H_{low})$ then $\pi(l_{i+1})$ is independent of τ_i and $\Pr[\pi(l_{i+1}) < r] \ge r$ (note that it is equal to r unless $l_{i+1} \in \text{TwoCC}$, in which case it is greater than r). Otherwise, let C be the connected component of H_{low} that contains l_{i+1} . For every $l \in V(C) \setminus \{l_{i+1}\}$, define A_l to be [0,1] if l is not in the condition τ_i , and otherwise $A_l := [0,r]$ if $\pi(l) < r$ and $A_l := [r,1]$ if $\pi(l) \ge r$. The condition τ_i contains this information. We apply Corollary 42 with $K = \{l_{i+1}\}$, $I = V(C) \setminus K$, and $A_{l_{i+1}} := [0,r]$, and get

$$\begin{aligned} \Pr[\pi(l_{i+1}) < r \mid \tau_i] &\geq r + \frac{\epsilon \gamma(r) \sum_{l:\{l_{i+1},l\} \in E(C)} T_v^-}{1 - \frac{2}{25} \epsilon(|E(C)| - 1)} \\ &\geq r + \frac{2 \epsilon \gamma(r) \frac{-\gamma(r)}{1 - r}}{1 - \frac{2}{25} \epsilon(|E(C)| - 1)} \\ &\geq r - 2.4 \epsilon \frac{\gamma^2(r)}{1 - r} . \end{aligned}$$

where the second inequality holds since T_v is either 0 or $\gamma(r)/r$ or $-\gamma(r)/(1-r)$, and l_{i+1} has at most two neighbors in C. The third inequality holds provided that $|C| \leq 17$ and $\epsilon \leq 0.13$. This means we can disentangle D by pessimistically sampling π such that $\Pr[\pi(l_i) < r] = r - \frac{2.4 \epsilon \gamma^2(r)}{1-r}$ for all $i \geq 1$, independently. Formally, define $\tilde{\gamma}(r) := \frac{2.4 \gamma^2(r)}{1-r}$ and let D' be the distribution under which $\pi(l_i) \sim D_{-\epsilon}^{\tilde{\gamma}}$ for every label $l_i \neq x$, and $\pi(x) \sim D_{\epsilon}^{\gamma_{\text{TwoCC}}}$. Then

$$\Pr_D[T_x] \ge \Pr_{D'}[T_x] \; .$$

We state the next theorem in slightly more general terms since we are going to re-use it in the next chapter.

Theorem 65. Let $\gamma_A, \gamma_{\text{rest}} : [0,1] \to \mathbb{R}_0^+$ be functions in the spirit of Definition 38 and assume that $\gamma_{\text{rest}}(r) = 0$ for all $r \ge 1/2$. ⁸ Let T_x be the critical clause tree for $x \in \text{TwoCC}$ as in (24). Let D' be a distribution on placements that samples $\pi(l) \sim D_{-\epsilon}^{\gamma_{\text{rest}}}$, i.e., $\Pr[\pi(l) < r] \ge r - \epsilon \gamma_{\text{rest}}(r)$ for every label l in T_x except x, and $\pi(x) \sim D_{\epsilon}^{\gamma_A}$. Then

$$\Pr_{D}[\operatorname{Cut}_{r}(T_{x})] \ge Q_{r} \cdot B(r) - \frac{2 \epsilon r \gamma_{\operatorname{rest}}(r) B(r)}{(1-r)^{3}} - o_{h}(1)$$
(25)

⁸Think of $\gamma_A = \gamma_{\text{TwoCC}}$ and $\gamma_{\text{rest}} = \tilde{\gamma}$.

$$\Pr_{D}[\operatorname{Cut}(T_{x})] \geq \operatorname{Cut2CC} - \epsilon(\operatorname{DFS2CC} + \operatorname{DFD2CC}) - \epsilon^{2} \operatorname{JUNK2CC} - o(1) .$$

where

$$\begin{split} \text{CUT2CC} &:= \int_0^1 Q_r B(r) \, dr = s_3 + \frac{104}{3} - 50 \, \ln(2) \approx s_3 + 0.009307 \\ \text{DFS2CC} &:= -\int_0^1 Q_r B(r) \phi_A(r) \, dr \ , \qquad \text{(damage from selfless bias of } x) \\ \text{DFD2CC} &:= \int_0^1 \frac{2r \gamma_{\text{rest}}(r) B(r)}{(1-r)^3} \, dr \ , \qquad \text{(damage from descendants' selfish bias)} \\ \text{JUNK2CC} &:= \int_0^1 \frac{2r \gamma_{\text{rest}}(r) B(r) \phi_A(r)}{(1-r)^3} \, dr \ . \end{split}$$

Note that CUT2CC does not depend on γ_{rest} and γ_A , so it is simply a universal constant, which is why we explicitly write it down.

Proof. For $r \geq 1/2$, the function $\gamma_{\text{rest}}(r)$ vanishes and thus $\Pr_D[\pi(l) < r] = r$, independently, for all labels l in the tree. Therefore, $\Pr_D[\operatorname{Cut}_r(T_x)] = 1 - o_h(1)$, and the claimed bound holds. So let us assume that r < 1/2. Since $\pi(a) \sim D_{-\epsilon}^{\gamma_{\text{rest}}}$ for all non-root labels, it holds that $\Pr[\pi(a) < r] = r - \epsilon \gamma_{\text{rest}}(r)$. One checks that B(r) is monotonically decreasing in r and therefore $B(r - \epsilon \gamma(r)) \geq B(r)$. Furthermore, Q_r is convex on [0, 1/2] and thus $Q_{r-\epsilon\gamma(r)} \geq Q_r - \epsilon\gamma(r)Q'_r = Q_r - \epsilon\gamma(r)\frac{2r}{(1-r)^3}$. Altogether,

$$Q_{r-\epsilon\gamma(r)} \cdot B(r-\epsilon\gamma(r)) \ge \left(Q_r - \frac{2\epsilon r\gamma(r)}{(1-r)^3}\right) \cdot B(r) .$$

= $Q_r \cdot B(r) - \frac{2\epsilon r\gamma(r)B(r)}{(1-r)^3}$

The claimed lower bound on $\Pr_{D'}[\operatorname{Cut}(T_x)]$ follows from taking the bound in (25) and taking the expectation over $\pi(x)$, using the fact that $\pi(x)$ has probability density $1 + \epsilon \phi(r)$ at r.

We apply the theorem with $\gamma_A(r) = \gamma_{\text{TwoCC}}(r) = 40r^{7/2}(1-2r)^2$ and $\gamma_{\text{rest}}(r) = \tilde{\gamma}(r) = \frac{2.4\gamma^2(r)}{1-r}$. Under this choice, we get

$$\begin{split} \mathrm{DFS2CC} &\leq 0.0455\\ \mathrm{DFD2CC} &\leq 0.0095\\ \mathrm{JUNK2CC} &\leq -0.019 \;. \end{split}$$

Plugging in these numbers into the formula in the theorem, we obtain:

Lemma 66. $\Pr_D[T_x] \ge s_3 + 0.009307 - 0.055 \epsilon$. for sufficiently large h.

7.8 The success probability of PPSZ under D

Let us collect the facts established in the previous sections.

- 1. $\frac{18}{17} |H_{\text{low}}| + 2 |H_{\text{high}}| + 2 |\text{TwoCC}| \ge |H|$, by (11).
- 2. $\operatorname{KL}(D||U) \leq 0.0064 \epsilon^2 |H_{\text{low}}| + \frac{5}{48 \ln(2)} f_{\text{KL}}(\epsilon) |\text{TwoCC}|$, by Lemma 45.

and

Next, for each variable x, let $(x \vee \overline{y} \vee \overline{z})$ be its canonical critical clause. Under D, the following hold:

- 3. $\Pr[\operatorname{Cut}(T_x)] \ge s_3 1.1 \epsilon \operatorname{Thr} + 0.001687 \epsilon$ when $\{y, z\} \in H_{\text{low}}$, by Lemma 55;
- 4. $\Pr[\operatorname{Cut}(T_x)] \ge s_3 1.1 \epsilon \operatorname{THR} + 0.9 \operatorname{THR}$ when $\{y, z\} \in H_{\text{high}}$, by Lemma 58;
- 5. $\Pr[\operatorname{Cut}(T_x)] \ge s_3 + 0.009307 0.055 \epsilon$ when $x \in \operatorname{TwoCC}$, by Lemma 66;
- 6. $\Pr[\operatorname{Cut}(T_x)] \ge s_3 1.1 \,\epsilon \text{THR}$ for all other x, by Lemma 53.

$$\Pr[\operatorname{ppsz}(F) \text{ succeeds}] = \mathop{\mathbb{E}}_{\pi} \left[2^{-n + \operatorname{Forced}(\pi)} \right] \qquad (by (1))$$
$$\geq 2^{-n + \mathop{\mathbb{E}}_{\pi \sim D}[\operatorname{Forced}(\pi)] - \operatorname{KL}(D||U)} \qquad (by (2))$$
$$= 2^{-n + s_3 n + gain},$$

where

$$\begin{aligned} \text{gain} &:= \ 0.001687 \, \epsilon |H_{\text{low}}| + 0.9 \, \text{Thr} |H_{\text{high}}| + (0.009307 - 0.055 \, \epsilon) |\text{TwoCC}| \\ &- 1.1 \, \epsilon \text{Thr} \, n \\ &- 0.0064 \, \epsilon^2 |H_{\text{low}}| - \frac{5}{48 \, \ln(2)} f_{\text{KL}}(\epsilon) |\text{TwoCC}| \\ &= \ (0.001687 \, \epsilon - 0.006404 \, \epsilon^2) |H_{\text{low}}| \\ &+ 0.9 \, \text{Thr} |H_{\text{high}}| \\ &+ (0.009307 - 0.055 \, \epsilon - 0.1503 \, f_{\text{KL}}(\epsilon)) |\text{TwoCC}| \\ &- 1.1 \, \epsilon \text{Thr} \, n \end{aligned}$$

Setting $\epsilon = 0.1$, the gain is at least

$$\begin{aligned} \text{gain} &\geq \frac{|H_{\text{low}}|}{9555} + 0.9 \text{ THR}|H_{\text{high}}| + \frac{|\text{TwoCC}|}{335} - 0.11 \text{ THR } n \\ &\geq \frac{\frac{18}{17}|H_{\text{low}}| + 0.9 \cdot 10118 \text{ THR}|H_{\text{high}}| + \frac{10118}{335}|\text{TwoCC}|}{10118} - 0.11 \text{ THR } n \end{aligned}$$

We set THR := $\frac{2}{0.9 \cdot 10118} \le \frac{1}{4553}$ so $0.9 \cdot 10118 \cdot \text{THR} \ge 2$ and

$$gain \ge \frac{\frac{18}{17}|H_{\text{low}}| + 2|H_{\text{high}}| + 3|\text{TwoCC}|}{10118} - \frac{n}{41391}$$
$$\ge \frac{|H|}{10118} - \frac{n}{41391} ,$$

where the last inequality follows from (11). This completes the proof of Theorem 35.

8 If there are many low-degree variables

As before TwoCC be the set of all variables that have two or more critical clauses. Additionally, let ID_i be the set of variables in $V \setminus TwoCC$ that have in-degree i in the critical clause graph. **Theorem** (Theorem 36, restated). The success probability of PPSZ is at least $2^{-n+s_3n+gain_2-o(n)}$

$$\operatorname{gain}_2 \ge \frac{|\mathrm{ID}_1| + 2 |\mathrm{ID}_0|}{1380}$$
,

where ID_i is the set of variables with in-degree *i* in the critical clause graph.

Proof. Let $\gamma_{\text{ID}_{0,1}} : [0,1] \to \mathbb{R}_0^+$ be a function in the spirit of Definition 38 and $\phi_{\text{ID}_{0,1}}$ its derivative. The idea is to sample

$$\pi(x) \sim D_{-\epsilon}^{\gamma_{\mathrm{ID}_{0,1}}}$$

i.e., with density $f_{\pi(x)}(r) = 1 - \epsilon \phi_{\text{ID}_{0,1}}(r)$ whenever $x \in \text{ID}_{0,1} := \text{ID}_0 \cup \text{ID}_1$. This means that x tends to come *later* in π ; indeed, $\Pr[\pi(x) < r] = r - \epsilon \gamma_{\text{ID}_{0,1}}(r)$. This is good for $\Pr[\text{Cut}(T_x)]$ but bad for every critical clause tree T_a that contains x. We will offset this damage by sampling $\pi(y)$ biased towards smaller values whenever (y, x) is an arc in the critical clause graph, i.e., if y is a "parent" of x. To make this formal, we choose another function $\gamma_{\text{PID}_{0,1}}$, the $\text{PID}_{0,1}$ standing for *parent* of ID_{0,1} and sample

$$\pi(y) \sim D_{\epsilon}^{\gamma_{\mathrm{pID}_{0,2}}}$$

whenever y is the parent of some $x \in ID_{0,1}$. To complicate things, a variable y can have no, one, or two children in $ID_{0,1}$; it could itself be in $ID_{0,1}$. Thus, we define a function γ_x for every variable x separately. To this end, let I_x be the indicator variable that is 1 if $x \in ID_{0,1}$ and 0 otherwise. We also need a function $\gamma_{TwoCC}(r)$ to handle variables $x \in TwoCC$. This $\gamma_{TwoCC}(r)$ is not the $40r^{7/2}(1-2r)^2$ from the previous section.

Definition 67. For each variable x, define $\gamma_x : [0,1] \to \mathbb{R}$ by

$$\gamma_x(r) := \begin{cases} -\gamma_{\mathrm{ID}_{0,1}}(r)I_x + \gamma_{\mathrm{PID}_{0,1}}(r)(I_y + I_z) & \text{if } x \notin \mathrm{TwoCC} \text{ and } (x \vee \bar{y} \vee \bar{z}) \text{ is its critical clause} \\ \gamma_{\mathrm{TwoCC}}(r) & \text{if } x \in \mathrm{TwoCC} \end{cases}$$

Then let D be the distribution on placements that samples each variable x independently from $D_{\epsilon}^{\gamma_x}$ independently.

The functions $\gamma_{\text{ID}_{0,1}}(r)$, $\gamma_{\text{pID}_{0,1}}(r)$ are defined to be 0 for $r \ge 1/2$; for r < 1/2, they are defined by

$$\begin{split} \gamma_{\mathrm{ID}_{0,1}}(r) &:= 10 \, r^2 (1-2 \, r)^2 \\ \gamma_{\mathrm{PID}_{0,1}}(r) &:= \frac{61}{6} \, r^3 (1-2 \, r)^2 \, , \\ \gamma_{\mathrm{TwoCC}}(r) &:= 20 \, r^3 (1-2 \, r) \, . \end{split}$$

We denote their derivatives by $\phi_{\text{ID}_{0,1}}$, $\phi_{\text{pID}_{0,1}}$, and ϕ_{TwoCC} , respectively.

8.1 $\Pr[Forced(x)]$ if $x \notin TwoCC$

Roughly speaking, we will show that $\Pr_D[\operatorname{Cut}(T_x)]$ is minimized if T_x is a full binary tree; no TwoCC-variables occur in it; no variables besides possibly x, y, z are in $\operatorname{ID}_{0,1}$. Then $\Pr_D[\operatorname{Cut}(T_x)]$ might differ from $\Pr_U[\operatorname{Cut}(T_x)]$ (under the uniform distribution) due to three factors: first, if $x \in \operatorname{ID}_{0,1}$ then $\pi(x)$ tends to be large, giving more weight to $\Pr[\operatorname{Cut}_r(T_x)]$ for large r; the effect is beneficial, i.e., it increases the cut probability; second, if $y \in \operatorname{ID}_{0,1}$, then then $\pi(y)$ tends to be larger, which has a detrimental effect on $\Pr_D[\operatorname{Cut}(T_x)]$; third, additionally if $y \in \operatorname{ID}_{0,1}$, then $\pi(x)$ gets a certain selfless bias, i.e., tends to be smaller (by Definition of γ_x in Definition 67), which further decreases $\Pr_D[\operatorname{Cut}(T_x)]$. The second and third effect of course also exist when $z \in \operatorname{ID}_{0,1}$. These effects are "almost" additive, and motivate the following definitions:

Definition 68. The following expressions quantify the benefits and damages from the biases as described above. The concrete numbers hold for our concrete choices of $\gamma_{\text{ID}_{0,1}}$ and $\gamma_{\text{PID}_{0,1}}$.

$$\begin{split} \text{BFS} &:= -\int_{0}^{1} \phi_{\text{ID}_{0,1}}(r) Q_r \, dr \;, \qquad (\text{``benefit from selfish bias''}) \\ &= 380 \ln(2) - \frac{790}{3} \ge 0.06259 \\ \text{DFC} &:= \int_{0}^{1} \gamma_{\text{ID}_{0,1}}(r) P_r(1 - Q_r) \, dr \;, \qquad (\text{``damage from children's selfish bias''}) \\ &= \frac{915}{4} - 330 \ln(2) \le 0.01144 \\ \text{DFS} &:= -\int_{0}^{1} \phi_{\text{PID}_{0,1}}(r) Q_r \, dr \;, \qquad (\text{``damage from selfless bias''}) \\ &= \frac{1586 \ln(2)}{3} - \frac{52765}{144} \le 0.0202 \\ \text{DFB} &:= \text{DFC} + \text{DFS} \qquad (\text{``damage from both''}) \\ &= \frac{596 \ln(2)}{3} - \frac{19825}{144} \le 0.03163 \;. \end{split}$$

Furthermore, we define

$$\begin{aligned} \operatorname{JUNK}_{1} &:= \max\left(0, -\int_{0}^{1}\phi_{\operatorname{ID}_{0,1}}(r)\gamma_{\operatorname{ID}_{0,1}}(r)P_{r}(1-Q_{r})\,dr\right) \\ &= 46800\,\ln(2) - \frac{227075}{7} \leq 0.00235 \\ \operatorname{JUNK}_{2} &:= \max\left(0, \int_{0}^{1}\phi_{\operatorname{PID}_{0,1}}(r)\gamma_{\operatorname{ID}_{0,1}}(r)P_{r}(1-Q_{r})\,dr\right) \\ &\quad \frac{8767591}{192} - 65880\,\ln(2) \leq 0.000184 \\ \operatorname{JUNK}_{1} &:= \operatorname{JUNK}_{1} + 2\,\operatorname{JUNK}_{2} \,. \end{aligned}$$

These terms quantify the extent to which the three described effects fail to be additive.

Theorem 69. Let $x \notin \text{TwoCC}$ and let $(x \lor \overline{y} \lor \overline{z})$ be its critical clause. Then

$$\Pr_{D}[\operatorname{Cut}(T_{x})] \ge s_{3} + \epsilon I_{x} \operatorname{BFS} - \epsilon (I_{y} + I_{z}) \operatorname{DFB} - \epsilon^{2} I_{x} (I_{y} + I_{z}) \operatorname{JUNK}_{1} - \epsilon^{2} (I_{y} + I_{z})^{2} \operatorname{JUNK}_{2} - o(1)$$

The impatient reader may skip the proof and continue directly in Section 8.2.

Proof of Theorem 69. Consider a critical clause tree T_x . Several variable labels in it might have a selfish bias (they tend to have high π -values) or a selfless bias (they tend to have low π -values). The next lemma argues that these biases cancel out, except for any biases of the root and its two children. To put it differently, in the worst case, no variables occurring in T_x are in TwoCC, and no variables beyond possibly the root and its children are in ID_{0,1}. To state the lemma formally, we need to define a distribution D^B on placements. **Definition 70.** Let T^{∞} be the infinite binary tree, and $B \subseteq V(T^{\infty})$. For each node $u \in V(T^{\infty})$, define

$$\gamma_u^B(r) := -\gamma_{\mathrm{ID}_{0,1}}(r)[u \in B] + \gamma_{\mathrm{pID}_{0,1}}(r)([v \in B] + [w \in B]) ,$$

where v, w are the two children of u in T^{∞} . Let D^B be the distribution that samples $\pi(u) \sim D_{\epsilon}^{\gamma_u}$ independently for every node.

Lemma 71. Let $x \in V \setminus \text{TwoCC}$ and let $(x \vee \overline{y} \vee \overline{z})$ be its unique critical clause. Let T^{∞} be the infinite binary tree, and let $\operatorname{root}(T^{\infty})$, L, R denote the root of T^{∞} , its left child, and its right child. Define $B \subseteq V(T^{\infty})$ by adding $\operatorname{root}(T^{\infty})$, L, and R to B if $x \in \text{ID}_{0,1}$, $y \in \text{ID}_{0,1}$, and $z \in \text{ID}_{0,1}$, respectively. Then $\Pr_D[\operatorname{Cut}(T_x)] \ge \Pr_{D^B}[T^{\infty}] - o(1)$.

Proof. First, we fix $r \in [0,1]$ and bound $\Pr[\operatorname{Cut}_r(T_x)]$ from below. If $r \geq 1/2$ then $\Pr[\pi(l) < r] \geq r$ for all labels (in fact with equality except for TwoCC-variablees). Thus, $\Pr[\operatorname{Cut}_r(T_x)] = 1 - o(1)$ if $r \geq 1/2$, and we are done.

From now on, fix some r < 1/2. We take T_x through a sequence of transformation steps, each of which does not increase $\Pr[\operatorname{Cut}_r]$ (or only by o(1)). Let T be T_x in the beginning. First, in a procedure similar to (but easier than) CLEANUP in the previous section, we can assign each node v of T a fresh label with the same distribution, i.e., $\gamma_v := \gamma_{\operatorname{varlabel}(v)}$.⁹ This does not increase $\Pr[\operatorname{Cut}(T)]$ at all (this can actually also be seen as a direct application of Lemma 7 from [8]).

Again as in the previous section, we can "normalize" T from level h' downwards: we make sure every node u of depth at least h' has two children (so T becomes infinite) and $\pi(u)$ is uniform over [0, 1]. This increases $\Pr[\text{Cut}]$ only by o(1).

Once T has this infinite shape, we can clean up TwoCC-variables and nodes with only one child. This is similar to the proofs in Section 7.3.3. As in that section, we denote by CLEANSUBTREE(T', u) the tree obtained from T' by replacing T_u with a copy of T^{∞} with fresh labels. Note that $\pi(l)$ is independent and uniform over [0, 1] for each fresh label.

Lemma 72 (TwoCC-cleanup in the irregular case). Let T be the (by now) infinite tree described above. Suppose u is a node in T with $\gamma_u = \gamma_{\text{TwoCC}}$. Then $\Pr[\text{Cut}_r(T)] \geq \Pr[\text{Cut}_r(\text{CLEANSUBTREE}(u,T)), \text{ provided that } \gamma_{\text{TwoCC}}(r) \geq \frac{2r\gamma_{\text{ID}_{0,1}}}{1-2r}.$

This is proved as Lemma D.1 in the appendix. The condition in the lemma justifies our particular choice for γ_{TwoCC} : the above inequality is satisfied with equality.

Lemma 73. Let T be the (by now) infinite tree described above. Suppose u is a node in T with an only child v. Then $\Pr[\operatorname{Cut}_r(T)] \ge \Pr[\operatorname{Cut}_r(\operatorname{CLEANSUBTREE}(u,T)))$, provided that $\epsilon \le \frac{4}{5}$.

This follows directly from Proposition 52 in the previous section. We just have to check the parameters: in this section, $\delta_{\max} = \epsilon \gamma_{\text{ID}_{0,1}}(r)$, and a simple calculation shows that the condition $r(1-2r) \ge 2 \delta_{\max}$, required by Proposition 52, holds.

We apply the operation CLEANSUBTREE(u, T) whenever applicable. In the resulting tree T, every node has two children (it is a T^{∞}), and $\gamma_v \neq 0$ only for "old" nodes v that already existed in T_x and are canonical. We define a subset $B \subseteq V(T)$ as follows: for each node $u \in V(T)$, add u to B if (1) it is an old node, i.e., already exists in T_x and has not been replaced in the above procedure; and (2) its old variable label $a := \text{varlabel}_{T_x}(u)$ is

⁹Once all labels are distinct, there is no need to distinguish between labels and nodes; we simply assume that v is the label of v, and thus can write γ_v instead of $\gamma_{\text{varlabel}(v)}$.

in ID_{0,1}. We will now compare γ_u , the "bias" of node *u* inherited from T_x through the transformations, to γ_u^B , the bias according to Definition 70.

Proposition 74. $\gamma_u(r) \ge \gamma_u^B(r)$.

Proof. Let v and w be the left and right child of u in T. If u is a new node, created during the cleanup process above, then v and w are new nodes, too, and both sides of the inequality are 0. So assume from now on that u is an old node, and let $a := \text{varlabel}_{T_x}(u)$. Note that a is canonical since otherwise the node u (and all its descendants) would have been replaced by new nodes, using Lemma 72 and Lemma 73. So a has a unique critical clause $(a \lor \overline{b} \lor \overline{c})$ and the children of u in T_x have labels b (left child) and c (right child). We see that

$$\gamma_u = -\gamma_{\mathrm{ID}_{0,1}}(r)[a \in \mathrm{ID}_{0,1}] + \gamma_{\mathrm{pID}_{0,1}}(r) \left([b \in \mathrm{ID}_{0,1}] + [c \in \mathrm{ID}_{0,1}]\right)$$

$$\gamma_u^B = -\gamma_{\mathrm{ID}_{0,1}}(r)[u \in B] + \gamma_{\mathrm{pID}_{0,1}}(r) \left([v \in B] + [w \in B]\right) .$$

We already assume that u is an old node, and thus $u \in B$ if and only if $a \in ID_{0,1}$, meaning the indicator variables $[a \in ID_{0,1}]$ and $[u \in B]$ are equal. Next, we will show that $[v \in B] \leq [b \in ID_{0,1}]$ and similar for w and c.

Indeed, if v is a new node, then trivially $0 = [v \in B] \leq [b \in ID_{0,1}]$. Otherwise, what is its old label? We already observed that u is a canonical node in T_x and its left child in T_x has label b. So variabel $T_x(v) = b$, and $[v \in B] = [b \in ID_{0,1}]$. The same holds for w and c. Since all γ 's are non-negative for all $r \in [0, 1]$, this proves the proposition. \Box

Note that we need the fact that u is canonical; otherwise, we could not guarantee any relation between the label of its children in T_x and the variables in its critical clause. This is why we have to handle TwoCC-variables and nodes with an only child separately and need to go through the pains of proving Lemma 72 and Lemma 73.

We conclude that $\Pr_D[\pi(u) < r] \ge \Pr_{D^B}[\pi(u) < r]$ and therefore can bound $\Pr_D[\operatorname{Cut}_r(T)] \ge \Pr_{D^B}[\operatorname{Cut}_r(T)]$. This is good since $\Pr_{D^B}[\operatorname{Cut}_r(T)]$ can be evaluated purely in terms of B, and we can completely forget about the original T_x and its labels. In what follows, we will show that we can reduce the set B until it contains nothing except possibly the root of T and its children, while not increasing $\Pr_{D^B}[\operatorname{Cut}_r(T)]$.

Lemma 75. Let v be a node in B of maximum distance from the root. If that distance is at least 2, then

$$\Pr_{D^B}[\operatorname{Cut}_r(T)] \ge \Pr_{D^B \setminus \{v,w\}}[\operatorname{Cut}_r(T)] ,$$

where w is the sibling of v in T (the other child of the parent of u).

Proof. Let u be the parent of v and w the other child of u. By assumption, the depth of v is at least two, so u is not the root. For brevity, we write $D_1 := D^B$ and $D_2 := D^{B \setminus \{v,w\}}$. Let T_u be the subtree of T rooted at u. It suffices to show that $\Pr_{D_1}[\operatorname{wCut}(T_u)] \geq \Pr_{D_2}[\operatorname{wCut}(T_u)]$. For $t \in \{u, v, w\}$ write $I_t := [t \in B]$. Note that

$$\begin{split} \gamma_v^B &= -\gamma_{\mathrm{ID}_{0,1}}(r)I_v & (\text{since the children of } v \text{ are not in } B) \\ \gamma_w^B &= -\gamma_{\mathrm{ID}_{0,1}}(r)I_w & (\text{since the children of } w \text{ are not in } B) \\ \gamma_u^B &= -\gamma_{\mathrm{ID}_{0,1}}(r)I_u + \gamma_{\mathrm{PID}_{0,1}}(r)(I_v + I_w) \\ \gamma_u^{B \setminus \{v,w\}} &= -\gamma_{\mathrm{ID}_{0,1}}(r)I_u \ . \end{split}$$

For $t \in \{u, v, w\}$ define δ_t so that $\Pr_{D_1}[\pi(t) < r] = r - \delta_t$; additionally define δ_u^{new} such that $\Pr_{D_2}[\pi(u) < r] = r - \delta_u^{\text{new}}$. Note that $\Pr_{D_2}[\pi(v) < r] = r$ and similar for w, so no need to define δ_v^{new} . We check that

$$\begin{split} \delta_u^{\text{new}} &= \epsilon \gamma_{\text{ID}_{0,1}}(r) I_u \\ \delta_v &= \epsilon \gamma_{\text{ID}_{0,1}}(r) I_v \\ \delta_w &= \epsilon \gamma_{\text{ID}_{0,1}}(r) I_w \\ \delta_u &= \delta_u^{\text{new}} - \epsilon \gamma_{\text{pID}_{0,1}}(r) (I_v + I_w) \;. \end{split}$$

With this notation, the "new" cut probability is

$$\Pr_{D_2}[\operatorname{wCut}_r(T_u)] = \Pr_{D_2}[\pi(u) < r \lor \operatorname{Cut}_r(T_u)]$$

$$= (r - \delta_u^{\operatorname{new}}) \lor Q_r \qquad \text{(since everything is unbiased below } u)$$

$$= (r - \delta_u^{\operatorname{new}}) + (1 - r + \delta_u^{\operatorname{new}})Q_r$$

$$= r - \delta_u^{\operatorname{new}} + (1 - r)Q_r + \delta_u^{\operatorname{new}}Q_r$$

$$= P_r - \delta_u^{\operatorname{new}}(1 - Q_r) . \qquad (26)$$

The "old" probability $\Pr_{D_1}[\operatorname{wCut}_r(T_u)]$ is a bit more tedious to compute. Analogous to (26), we obtain

$$\Pr_{D_1}[\operatorname{wCut}_r(T_v)] = P_r - \delta_v(1 - Q_r) ,$$

$$\Pr_{D_1}[\operatorname{wCut}_r(T_w)] = P_r - \delta_w(1 - Q_r) ,$$

and therefore

$$\begin{split} &\Pr_{D_1}[\mathrm{wCut}_r(T_u)] = (r - \delta_u) \lor \Pr_{D_1}[\mathrm{Cut}_r(T_u)] \\ &= (r - \delta_u) \lor \left(\Pr_{D_1}[\mathrm{wCut}_r(T_v)] \cdot \Pr_{D_1}[\mathrm{wCut}_r(T_w)]\right) \\ &= (r - \delta_u) \lor ((P_r - \delta_v(1 - Q_r)) \cdot (P_r - \delta_w(1 - Q_r))) \\ &= (r - \delta_u) \lor (Q_r - (\delta_v + \delta_w)(1 - Q_r)P_r + \delta_v\delta_w(1 - Q_r)^2) \\ &= r - \delta_u + (1 - r + \delta_u) \left(Q_r - (\delta_v + \delta_w)(1 - Q_r)P_r + \delta_v\delta_w(1 - Q_r)^2\right) \\ &= r - \delta_u + (1 - r)Q_r - (1 - r)(\delta_v + \delta_w)(1 - Q_r)P_r + \delta_uQ_r \\ &+ (1 - r + \delta_u)\delta_v\delta_w(1 - Q_r)^2 - \delta_u(\delta_v + \delta_w)(1 - Q_r)P_r \\ &= P_r - (1 - Q_r)(\delta_u + r\delta_v + r\delta_w) \qquad (\text{since } (1 - r)P_r = r) \\ &+ (1 - r + \delta_u)\delta_v\delta_w(1 - Q_r)^2 - \delta_u(\delta_v + \delta_w)P_r(1 - Q_r) \ . \end{split}$$

We compute the difference between "old" and "new":

$$\Pr_{D_{1}}[\operatorname{wCut}_{r}(T_{u})] - \Pr_{D_{2}}[\operatorname{wCut}_{r}(T_{u})] = (1 - Q_{r})(\delta_{u}^{\operatorname{new}} - \delta_{u} - r\delta_{v} - r\delta_{w}) + (1 - r + \delta_{u})\delta_{v}\delta_{w}(1 - Q_{r})^{2} - \delta_{u}(\delta_{v} + \delta_{w})P_{r}(1 - Q_{r}) = (1 - Q_{r})\epsilon(I_{v} + I_{w})(\gamma_{\operatorname{pID}_{0,1}}(r) - r\gamma_{\operatorname{ID}_{0,1}}(r))$$
(27)
+ $(1 - r + \delta_{u})\delta_{v}\delta_{w}(1 - Q_{r})^{2} - \delta_{u}(\delta_{v} + \delta_{w})P_{r}(1 - Q_{r})$ (28)

Since $\gamma_{\text{pID}_{0,1}}(r) \ge r\gamma_{\text{ID}_{0,1}}(r)$ for all r, the expression in (27) is non-negative. Morally, (27) is linear in ϵ while (28) is quadratic; so even if (28) is negative, it should not matter

too much. However, to rigorously show that the overall difference is non-negative, we distinguish a couple of cases.

Case 1. $u \notin B$. Then the second term of (28) vanishes since $\delta_u = 0$; the first term is non-negative: the factor $(1 - r + \delta_u)$ is a probability and is therefore non-negative; δ_v and δ_w are both at most 0, so $\delta_v \delta_w \ge 0$.

Case 2. $u \in B$ and $v, w \notin B$. Then (28) vanishes.

Case 3. $u, v, w \in B$. Then $\delta_v = \delta_w > 0$. Dividing by $1 - Q_r$ and δ_v , we see that (28) ≥ 0 if and only if

$$(1 - r + \delta_u)\delta_v(1 - Q_r) - 2\,\delta_u P_r \ge 0 \tag{29}$$

We assume that $I_u = I_v = I_w = 1$, and therefore $\delta_u = \epsilon \gamma_{\text{ID}_{0,1}} - 2 \epsilon \gamma_{\text{PID}_{0,1}}$. If δ_u is negative, then both summands of (29) are non-negative (we check that $1 - r + \delta_u$ is always positive for $r \leq 1/2$ and $\epsilon \leq 1$, by a wide margin). If δ_u is positive, then

$$(29) \ge (1-r)\delta_v(1-Q_r) - 2\delta_u P_r$$

= $(1-r)(1-Q_r)\epsilon\gamma_{\mathrm{ID}_{0,1}} - 2\epsilon P_r \left(\gamma_{\mathrm{ID}_{0,1}} - 2\gamma_{\mathrm{pID}_{0,1}}\right)$
 $\le (1-r)(1-Q_r)\epsilon\gamma_{\mathrm{ID}_{0,1}} - 2\epsilon P_r \left(\gamma_{\mathrm{ID}_{0,1}} - 2r\gamma_{\mathrm{ID}_{0,1}}\right)$

where the last inequality follows since $\gamma_{\text{pID}_{0,1}}(r) = \frac{61}{60} \cdot r \cdot \gamma_{\text{ID}_{0,1}}(r) \ge r \cdot \gamma_{\text{ID}_{0,1}}(r)$. We divide this by $\epsilon \gamma_{\text{ID}_{0,1}}(r)$, which is positive, and continue:

$$\begin{aligned} (1-r)(1-Q_r) - 2P_r(1-2r) &= \frac{(1-r)(1-2r)}{(1-r)^2} - \frac{2r(1-2r)}{1-r} \\ &= \frac{(1-2r) - 2r(1-2r)}{1-r} \\ &= \frac{(1-2r)^2}{1-r} > 0 \;. \end{aligned}$$

This shows that $(29) \ge 0$ and therefore $(28) \ge 0$.

Case 4. $I_u = 1$ and exactly one of $I_v, I_w = 1$. Say $I_v = 1$ and $I_w = 0$. The difference is

$$(27) + (28) = (1 - Q_r)\epsilon(\gamma_{\text{PID}_{0,1}}(r) - r\gamma_{\text{ID}_{0,1}}(r)) - \delta_u \delta_v P_r(1 - Q_r) = (1 - Q_r)\epsilon(\gamma_{\text{PID}_{0,1}}(r) - r\gamma_{\text{ID}_{0,1}}(r)) - \epsilon^2 \gamma_{\text{ID}_{0,1}}(\gamma_{\text{ID}_{0,1}} - \gamma_{\text{PID}_{0,1}})P_r(1 - Q_r) .$$

We divide this by $(1 - Q_r)\epsilon\gamma_{\text{ID}_{0,1}}$, which is positive, keeping in mind that $\frac{\gamma_{\text{PID}_{0,1}}(r)}{\gamma_{\text{ID}_{0,1}}} = \frac{61 r}{60}$, and the above becomes

$$\frac{r}{60} - \epsilon(\gamma_{\mathrm{ID}_{0,1}} - \gamma_{\mathrm{pID}_{0,1}})P_r$$

$$\geq \frac{r}{60} - \epsilon(\gamma_{\mathrm{ID}_{0,1}} - r\gamma_{\mathrm{ID}_{0,1}})P_r \qquad (\text{since } \gamma_{\mathrm{pID}_{0,1}}(r) \geq r\gamma_{\mathrm{ID}_{0,1}}(r))$$

$$= \frac{r}{60} - \epsilon r\gamma_{\mathrm{ID}_{0,1}} , \qquad (\text{since } P_r = \frac{r}{1-r})$$

which is non-negative if and only if $\epsilon \gamma_{\text{ID}_{0,1}}(r) \leq 60$. Since $\gamma_{\text{ID}_{0,1}}(r) = 10 r^2 (1-2r)^2 \leq \frac{10}{256}$ for all $0 \leq r \leq 1/2$, this holds for all $\epsilon \leq \frac{256}{600}$. This concludes the proof of Lemma 75. \Box

We apply Lemma 75 as long as B contains a node of depth at least 2. Once there is no such node anymore, we conclude that B does not contain any node outside the root and its children. This completes the proof of Lemma 71.

Informally, Lemma 71 states that biases cancel and we can pretend that $TwoCC = \emptyset$ and T_x contains no ID_{0,1}-variables except possibly x, y, z. After applying Lemma 71, we are left analyzing $\Pr_{D^B}[T^{\infty}]$. Notation becomes simpler if we assume that the root, left child, and right child of T^{∞} have labels x, y, z, respectively. Then B is simply $\{x, y, z\} \cap$ ID_{0,1}. Things are now simple enough to calculate $\Pr_{D^B}[\operatorname{Cut}(T^{\infty})]$. We abbreviate $\gamma_1 :=$ $\gamma_{\mathrm{ID}_{0,1}}$ and $\gamma_2 := \gamma_{\mathrm{pID}_{0,1}}$. With this notation,

$$\Pr[\pi(y) < r] = r - \epsilon \gamma_1(r) I_y =: r - \delta_y ,$$

$$\Pr[\pi(z) < r] = r - \epsilon \gamma_1(r) I_z =: r - \delta_z ,$$

$$\Pr[\pi(x) < r] = r - \epsilon \gamma_1(r) I_x + \epsilon \gamma_2(r) (I_y + I_z) =: r - \delta_x ,$$
(30)

and $\Pr[\pi(l) < r] = r$ for all other labels l; all probabilities are under D^B . Denote by T_u and T_z the subtrees of T^{∞} rooted at y (i.e., the left child of the root) and z, respectively.

$$\begin{aligned} \Pr_{D^B}[\operatorname{Cut}_r(T^\infty)] &= \Pr[\pi(y) < r \lor \operatorname{Cut}_r(T_y)] \cdot \Pr[\pi(z) < r \lor \operatorname{Cut}_r(T_z)] \\ &= \prod_{l \in \{y, z\}} (r - \delta_l + (1 - r + \delta_l)Q_r) \\ &= \prod_{l \in \{y, z\}} (r + (1 - r)Q_r - \delta_l(1 - Q_r)) \\ &= \prod_{l \in \{y, z\}} (P_r - \delta_l(1 - Q_r)) \\ &= Q_r - (\delta_y + \delta_z)P_r(1 - Q_r) + \delta_y \delta_z (1 - Q_r)^2 \\ &\geq Q_r - (I_y + I_z)\epsilon\gamma(r)P_r(1 - Q_r) \;, \end{aligned}$$

since $\delta_y, \delta_z \ge 0$. We get $\Pr_{D^B}[\operatorname{Cut}(T^{\infty})]$ from by taking the expectation of the above with $r = \pi(x)$; we have to be aware that $\pi(x)$ is not uniform if $I_x = 1$. That is, we integrate the above expression, multiplied with the density of $\pi(x)$. The density of $\pi(x)$ at r is $\gamma'_x(r)$, which can be obtained by differentiating (30):

$$1 - \epsilon \phi_1(r) I_x + \epsilon \phi_2(r) (I_y + I_z)$$

where $\phi_i(r) := \gamma'_i(r)$ for i = 1, 2. Writing $\bar{Q}_r := 1 - Q(r)$, we see that $\Pr_{D^B}[\operatorname{Cut}(T^{\infty})]$ is

$$\begin{split} &\int_{0}^{1} \left(Q_{r} - \epsilon \gamma_{1}(r) P_{r} \bar{Q}_{r}(I_{y} + I_{z}) \right) \left(1 - \epsilon \phi_{1}(r) I_{x} + \epsilon \phi_{2}(r) (I_{y} + I_{z}) \right) dr \\ &= s_{3} - \epsilon (I_{y} + I_{z}) \int_{0}^{1} \gamma_{1}(r) P_{r} \bar{Q}_{r} dr - \epsilon I_{x} \int_{0}^{1} \phi_{1}(r) Q_{r} dr + \epsilon (I_{y} + I_{z}) \int_{0}^{1} \phi_{2}(r) Q_{r} dr \\ &+ \epsilon^{2} I_{x}(I_{y} + I_{z}) \int_{0}^{1} \phi_{1}(r) \gamma_{1}(r) P_{r} \bar{Q}_{r} dr - \epsilon^{2} (I_{y} + I_{z})^{2} \int_{0}^{1} \phi_{2}(r) \gamma_{1}(r) P_{r} \bar{Q}_{r} dr \; . \end{split}$$

Using the notation from Definition 68, we can write

$$\Pr_{D^B}[\operatorname{Cut}(T^{\infty})] \ge s_3 + \epsilon I_x \operatorname{BFS} - \epsilon (I_y + I_z)(\operatorname{DFB}) - \epsilon^2 I_x (I_y + I_z) \operatorname{JUNK}_1 - \epsilon^2 (I_y + I_z)^2 \operatorname{JUNK}_2$$

This finishes the proof of Theorem 69.

This finishes the proof of Theorem 69.

8.2 Summing over all $x \in V \setminus \text{TwoCC}$

Next, we sum the expression in Theorem 69 over all $x \notin \text{TwoCC}$: for succinctness, we separately sum terms constant, linear, and quadratic in ϵ , respectively. Summing up $s_3 - o(1)$ gives at least $s_3(n - |\text{TwoCC}|) - o(n)$. The most interesting part is the one linear in ϵ . Recall that y, z are the other variables in x's canonical critical clause $(x \vee \overline{y} \vee \overline{z})$; that is, they are the variables for which (x, y), (x, z) are arcs in the critical clause graph. We write $x \to y, x \to z$ for succinctness.

$$\begin{split} &\sum_{x \in V \setminus \text{TwoCC}} \left(I_x \text{BFS} - \sum_{y:x \to y} I_y \text{DFB} \right) \\ = |\text{ID}_{0,1}| \cdot \text{BFS} - \sum_{x \in V \setminus \text{TwoCC}} \sum_{y:x \to y} I_y \text{DFB} \\ = |\text{ID}_{0,1}| \cdot \text{BFS} - \text{DFB} \sum_{y \in V} I_y \sum_{x \in V \setminus \text{TwoCC}} [x \to y] \\ = |\text{ID}_{0,1}| \cdot \text{BFS} - \text{DFB} \sum_{y \in \text{ID}_{0,1}} \sum_{x \in V \setminus \text{TwoCC}} [x \to y] \quad (\text{since } I_y = 0 \text{ if } y \notin \text{ID}_{0,1}) \\ \ge |\text{ID}_{0,1}| \cdot \text{BFS} - \text{DFB} \sum_{y \in \text{ID}_{0,1}} \deg_{\text{in}}(y) \quad (``\geq `` \text{ since some } x : x \to y \text{ might be in TwoCC}) \\ = |\text{ID}_{0,1}| \cdot \text{BFS} - |\text{ID}_1| \cdot \text{DFB} . \quad (\text{since } \deg_{\text{in}}(y) \text{ is 1 if } y \in \text{ID}_1 \text{ and 0 if } y \in \text{ID}_0) \end{split}$$

Next, let us sum up the terms with factor $\epsilon^2 JUNK_1$:

$$\sum_{\substack{x \in V \setminus \text{TwoCC} \\ C_x = : (x \lor \bar{y} \lor \bar{z})}} I_x(I_y + I_z) = \sum_{x \in V \setminus \text{TwoCC}} \sum_{y : x \to y} I_x I_y$$

= the number of arcs within $ID_{0,1}$ in the critical clause graph

$$\leq |\mathrm{ID}_1|$$

where the last inequality holds since every arc (x, y) within $ID_{0,1}$ has $y \in ID_1$, and by definition of ID_1 the variable y has at most one such incoming arc. Finally, the coefficient of $\epsilon^2 JUNK_2$ is

$$\sum_{\substack{x \in V \setminus \text{TwoCC} \\ C_x = : (x \lor \bar{y} \lor \bar{z})}} (I_y + I_z)^2 = \sum_{x \in V \setminus \text{TwoCC}} \left(\sum_{y:x \to y} I_y\right)^2$$

$$\leq \sum_{x \in V \setminus \text{TwoCC}} 2 \sum_{y:x \to y} I_y \qquad \text{(since the inner sum is at most 2)}$$

$$= 2 \sum_{y} \sum_{x \in V \setminus \text{TwoCC}} I_y$$

$$\leq 2 \sum_{y} \sum_{x \in V:x \to y} I_y = 2 \sum_{y \in \text{ID}_{0,1}} \deg_{\text{in}}(y) = 2 |\text{ID}_1| .$$

Altogether, we get

$$\sum_{x \notin \text{TwoCC}} \Pr_{D}[\text{Cut}(T_{x})] \geq s_{3}(n - |\text{TwoCC}|) - o(n) + \epsilon \left(|\text{ID}_{0,1}|\text{BFS} - |\text{ID}_{1}|\text{DFB}) - \epsilon^{2}|\text{ID}_{1}|\left(\text{JUNK}_{1} + 2 \text{JUNK}_{2}\right) \right)$$
(31)

8.3 $\Pr[Forced(x)]$ is large when x has two critical clauses.

Note that $\Pr[\pi(l) < r] = r - \epsilon \gamma_l(r) \ge r - \epsilon \gamma_{\text{ID}_{0,1}}(r)$ holds for every label l in T_x , and $\pi(x) \sim D_{\epsilon}^{\gamma_{\text{TwoCC}}}$. We can directly apply Theorem 65 with $\gamma_A = \gamma_{\text{TwoCC}}$ and $\gamma_{\text{rest}} = \gamma_{\text{ID}_{0,1}}$ and conclude that

$$\Pr_{D}[\operatorname{Cut}(T_{x})] \ge s_{3} + \operatorname{BONUS2CC} - \epsilon(\operatorname{DFS2CC} + \operatorname{DFD2CC}) - \epsilon^{2} \operatorname{JUNK2CC} - o(1) .$$
(32)

with

$$\begin{aligned} \text{Bonus2CC} &:= \int_0^1 Q_r B(r) \, dr - s_3 = \frac{104}{3} - 50 \, \ln(2) \approx 0.009307 \\ \text{DFS2CC} &:= -\int_0^1 Q_r B(r) \phi_{\text{TwoCC}}(r) \, dr \ , \\ &= \frac{39094}{3} - 18800 \, \ln(2) \leq 0.16634 \\ \text{DFD2CC} &:= \int_0^1 \frac{2r \gamma_{\text{ID}_{0,1}}(r) B(r)}{(1-r)^3} \, dr \ , \\ &= 11420 \, \ln(2) - \frac{23747}{3} \leq 0.074135 \\ \text{JUNK2CC} &:= \int_0^1 \frac{2r \gamma_{\text{ID}_{0,1}}(r) B(r) \phi_{\text{TwoCC}}(r)}{(1-r)^3} \, dr \\ &= \frac{17923400}{7} - 3694000 \, \ln(2) \approx 0.03125 \end{aligned}$$

8.4 Success Probability of PPSZ in terms of TwoCC and $ID_{0,1}$

The overall success probability of PPSZ can now be bounded as follows:

$$\log_{2} \Pr[\text{ppsz succeeds}] = \log_{2} \mathbb{E} \left[2^{-n + \sum_{x} \Pr[\text{Forced}(x,\pi)]} \right]$$
$$\geq \log_{2} \mathbb{E} \left[2^{-n + \sum_{x} \Pr[\text{Cut}(T_{x})]} \right]$$
$$\geq -n + \sum_{x} \Pr_{D}[\text{Cut}(T_{x})] - \text{KL}(D||U)$$
(33)

Since D acts independently on all variables, computing KL(D||U) should be straightforward; still, it causes a slight headache since variables come in seven "flavors", meaning γ_x is chosen among seven possible functions. To be more precise, j = 0, 1, 2 we define

 $A_{1,j} := \{ x \in \mathrm{ID}_{0,1} \mid \text{exactly } j \text{ of } x \text{'s children are in } \mathrm{ID}_{0,1} \}$ $A_{0,j} := \{ x \in V \setminus (\mathrm{ID}_{0,1} \cup \mathrm{TwoCC}) \mid \text{exactly } j \text{ of } x \text{'s children are in } \mathrm{ID}_{0,1} \}$

Looking back at Definition 67, we see that $\gamma_x = -i\gamma_{\text{ID}_{0,1}} + j \cdot \gamma_{\text{pID}_{0,1}} =: \gamma_{i,j}$ when $x \in A_{i,j}$ In particular, $\gamma_x = 0$ for $x \in B_{0,0}$, so $\pi(x)$ is uniform, and these variables contribute nothing to KL. We define $\phi_{i,j}(r) := \gamma'_{i,j}(r) dr$, $\Psi_{i,j} := \int_0^1 \phi_{i,j}^2(r)$, and $D_{i,j} := D_{\epsilon}^{\gamma_{i,j}}$. With this notation, we can write

$$\mathrm{KL}(D||U) = \sum_{i=0}^{1} \sum_{j=0}^{2} \mathrm{KL}(D_{i,j}||U) \cdot |A_{i,j}| + \mathrm{KL}(D_{\epsilon}^{\gamma_{\mathrm{TwoCC}}}||U) \cdot |\mathrm{TwoCC}|$$

Using Lemma 43 we obtain that

$$\operatorname{KL}(D_{i,j}) \le \frac{\Psi_{i,j}}{\ln(2)} \cdot f_{\operatorname{KL}}(\epsilon)$$

for $f_{\mathrm{KL}}(\epsilon) = (1-\epsilon) \ln(1-\epsilon) + \epsilon$. One checks that $\Psi_{1,2} \leq \Psi_{1,1} \leq \Psi_{1,0}$ for our choice of $\gamma_{\mathrm{ID}_{0,1}}$ and $\gamma_{\mathrm{pID}_{0,1}}$. Therefore,

$$\sum_{j=0}^{2} \operatorname{KL}(D_{A_{1,j}}||U) \cdot |A_{1,j}| \leq \frac{f_{\operatorname{KL}}(\epsilon)}{\ln(2)} \left(\Psi_{A_{1,0}}|A_{1,0}| + \Psi_{A_{1,1}}|A_{1,1}| + \Psi_{A_{1,2}}|A_{1,2}| \right)$$

$$\leq \frac{f_{\operatorname{KL}}(\epsilon)}{\ln(2)} \cdot \Psi_{A_{1,0}}|A_{1,0} \cup A_{1,1} \cup A_{1,2}| = \frac{f_{\operatorname{KL}}(\epsilon)}{\ln(2)} \cdot \Psi_{A_{1,0}}|\operatorname{ID}_{0,1}| .$$
(34)

For $A_{0,j}$, note that

$$\Psi_{A_{0,j}} = \int_0^1 \phi_{A_{0,j}}^2(r) \, dr = j^2 \Psi_{A_{0,j}}$$

and therefore

$$\sum_{j=0}^{2} \operatorname{KL}(D_{A_{0,j}}||U) \cdot |A_{1,j}| \leq \frac{f_{\operatorname{KL}}(\epsilon)}{\ln(2)} \sum_{j=0}^{2} \Psi_{A_{0,j}}$$

$$= \frac{f_{\operatorname{KL}}(\epsilon)}{\ln(2)} \Psi_{A_{0,2}} \left(\frac{1}{4}|A_{0,1}| + |A_{0,2}|\right)$$

$$\leq \frac{f_{\operatorname{KL}}(\epsilon)}{2\ln(2)} \Psi_{A_{0,2}} \left(|A_{0,1}| + 2|A_{0,2}|\right)$$

$$\leq \frac{f_{\operatorname{KL}}(\epsilon)}{2\ln(2)} \Psi_{A_{0,2}} |\operatorname{ID}_{1}| . \tag{35}$$

To see why the last inequality holds, note that $A_{0,j}$ is, by definition, the set of $x \in V \setminus \text{TwoCC}$ that are not in $\text{ID}_{0,1}$ but have j arcs into $\text{ID}_{0,1}$. Thus, the expression $|A_{0,1}| + 2|A_{0,2}|$ is at most the number of arcs in the critical clause graph going from outside $\text{ID}_{0,1}$ into $\text{ID}_{0,1}$ (at most since some arcs might come from TwoCC-variables). This number is at most $|\text{ID}_1|$. To determine $\text{KL}(D_{\epsilon}^{\gamma_{\text{TwoCC}}}||U)$, we cannot apply Lemma 43 directly since it requires $|\phi| \leq 1$. In fact, $\phi_{\text{TwoCC}}(r)$ is -5 for r = 1/2, and this is its maximal absolute value. Thus, $|\frac{1}{5}\phi_{\text{TwoCC}}(r)| \leq 1$, and therefore

$$\operatorname{KL}(D_{\epsilon}^{\gamma_{\operatorname{TwoCC}}}||U) = \operatorname{KL}(D_{5\epsilon}^{\gamma_{\operatorname{TwoCC}}/5}||U) \le \frac{f_{\operatorname{KL}}(5\epsilon)}{\ln(2)} \frac{1}{25} \Psi_{\operatorname{TwoCC}}$$
(36)

for $\Psi_{\text{TwoCC}} := \int_0^1 \phi_{\text{TwoCC}}^2(r) \, dr = \frac{15}{14}$. Adding (34), (35), and (36) gives

$$\begin{aligned}
\mathrm{KL}(D||U) &\leq \frac{f_{\mathrm{KL}}(\epsilon)}{\ln(2)} \left(\Psi_{1,0}|\mathrm{ID}_{0,1}| + \frac{\Psi_{0,2}}{2}|\mathrm{ID}_{1}| \right) + \frac{f_{\mathrm{KL}}(5\,\epsilon)}{25\,\ln(2)} \Psi_{\mathrm{TwoCC}}|\mathrm{TwoCC}| \\
&= \frac{f_{\mathrm{KL}}(\epsilon)}{\ln(2)} \left(\frac{5}{21} |\mathrm{ID}_{0,1}| + \frac{3721}{90720} |\mathrm{ID}_{1}| \right) + \frac{f_{\mathrm{KL}}(5\,\epsilon)}{25\,\ln(2)} \frac{15}{14} |\mathrm{TwoCC}| \end{aligned} \tag{37}$$

This bound is more or less tight (up to the slack in Lemma 43: it might so happen that (1) no $ID_{0,1}$ -variable x has an out-neighbor $y \in ID_{0,1}$; (2) ID_1 -variables occur in "pairs", i.e., if (x, y), (x, z) the out-arcs of x, then $y \in ID_1$ if and only if $z \in ID_1$. In this scenario,

 $A_{1,1} = A_{1,2} = A_{0,1} = \emptyset$, $|A_{1,0}| = |ID_{0,1}|$, and $|A_{0,2}| = |ID_{0,1}|/2$. Combining (37), (31), and (33), we see that the "gain" $\log_2 \Pr[\text{success}] + n - s_3 n$ is at least

$$\begin{split} \epsilon(|\mathrm{ID}_{0,1}|\mathrm{BFS} - |\mathrm{ID}_{1}|\mathrm{DFB}) &- \epsilon^{2}|\mathrm{ID}_{1}|\mathrm{JUNK} - \frac{f_{\mathrm{KL}}(\epsilon)}{\ln(2)} \left(\Psi_{A_{1,0}}|\mathrm{ID}_{0,1}| + \frac{\Psi_{A_{0,2}}}{2}|\mathrm{ID}_{1}|\right) \\ &+ \left(\mathrm{BONUS2CC} - \epsilon(\mathrm{DFS2CC} + \mathrm{DFD2CC}) - \epsilon^{2}\mathrm{JUNK2CC} - \frac{f_{\mathrm{KL}}(5\epsilon)}{25\ln(2)}\frac{15}{14}\right)|\mathrm{TwoCC}| \\ &= |\mathrm{ID}_{1}| \left(\epsilon(\mathrm{BFS} - \mathrm{DFB}) - \epsilon^{2}\mathrm{JUNK} - \frac{\Psi_{A_{1,0}} + \Psi_{A_{0,2}}/2}{\ln(2)}f_{\mathrm{KL}}(\epsilon)\right) \\ &+ |\mathrm{ID}_{0}| \left(\epsilon\mathrm{BFS} - \frac{f_{\mathrm{KL}}(\epsilon)\Psi_{A_{1,0}}}{\ln(2)}\right) \\ &+ \left(\mathrm{BONUS2CC} - \epsilon(\mathrm{DFS2CC} + \mathrm{DFD2CC}) - \epsilon^{2}\mathrm{JUNK2CC} - \frac{f_{\mathrm{KL}}(5\epsilon)}{25\ln(2)}\frac{15}{14}\right)|\mathrm{TwoCC}| \\ &\geq |\mathrm{ID}_{1}| \left(0.030966\epsilon - 0.0028\epsilon^{2} - 0.4027f_{\mathrm{KL}}(\epsilon)\right) \\ &+ |\mathrm{ID}_{0}| \left(0.06259\epsilon - 0.344f_{\mathrm{KL}}(\epsilon)\right) \\ &+ |\mathrm{TwoCC}| \left(0.009307 - 0.2405\epsilon - 0.03125\epsilon^{2} - 0.06183f_{\mathrm{KL}}(5\epsilon)\right) \end{split}$$

For $\epsilon = 0.029$, this is at least

$$\frac{|\mathrm{ID}_1|}{1380} + \frac{|\mathrm{ID}_0|}{600} + \frac{|\mathrm{TwoCC}|}{617} \ge \frac{|\mathrm{ID}_1| + 2\,|\mathrm{ID}_0| + 2\,|\mathrm{TwoCC}|}{1380}$$

This completes the proof of Theorem 36.

Acknowledgments

I am grateful to Robin Moser and Timon Hertli, with whom I intensively discussed the PPSZ algorithm, and some ideas used in this paper came from our discussions. I am grateful to Navid Talebanfard for fruitful discussions and helpful comments on this work.

References

- S. Baumer and R. Schuler. Improving a probabilistic 3-sat algorithm by dynamic search and independent clause pairs. In E. Giunchiglia and A. Tacchella, editors, *Theory and Applications of Satisfiability Testing, 6th International Conference, SAT* 2003. Santa Margherita Ligure, Italy, May 5-8, 2003 Selected Revised Papers, volume 2919 of Lecture Notes in Computer Science, pages 150–161. Springer, 2003.
- [2] T. M. Chan and R. Williams. Deterministic app, orthogonal vectors, and more: Quickly derandomizing razborov-smolensky. In R. Krauthgamer, editor, Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016, pages 1246–1255. SIAM, 2016.
- [3] T. D. Hansen, H. Kaplan, O. Zamir, and U. Zwick. Faster k-sat algorithms using biased-ppsz. In M. Charikar and E. Cohen, editors, Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019, pages 578–589. ACM, 2019.

- [4] T. Hertli. 3-SAT faster and simpler—unique-SAT bounds for PPSZ hold in general. In 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science—FOCS 2011, pages 277–284. IEEE Computer Soc., Los Alamitos, CA, 2011.
- [5] T. Hertli. Breaking the PPSZ Barrier for Unique 3-SAT. In J. Esparza, P. Fraigniaud, T. Husfeldt, and E. Koutsoupias, editors, Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I, volume 8572 of Lecture Notes in Computer Science, pages 600-611. Springer, 2014.
- [6] T. Hofmeister, U. Schöning, R. Schuler, and O. Watanabe. A probabilistic 3-sat algorithm further improved. In H. Alt and A. Ferreira, editors, STACS 2002, 19th Annual Symposium on Theoretical Aspects of Computer Science, Antibes - Juan les Pins, France, March 14-16, 2002, Proceedings, volume 2285 of Lecture Notes in Computer Science, pages 192–202. Springer, 2002.
- [7] S. Liu. Chain, generalization of covering code, and deterministic algorithm for ksat. In I. Chatzigiannakis, C. Kaklamanis, D. Marx, and D. Sannella, editors, 45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic, volume 107 of LIPIcs, pages 88:1–88:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- [8] R. Paturi, P. Pudlák, M. E. Saks, and F. Zane. An improved exponential-time algorithm for k-SAT. J. ACM, 52(3):337–364 (electronic), 2005.
- [9] R. Paturi, P. Pudlák, and F. Zane. Satisfiability coding lemma. Chicago J. Theoret. Comput. Sci., pages Article 11, 19 pp. (electronic), 1999.
- [10] P. Pudlák, D. Scheder, and N. Talebanfard. Tighter hard instances for PPSZ. In 44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland, pages 85:1–85:13, 2017.
- [11] T. Qin and O. Watanabe. An improvement of the algorithm of Hertli for the unique 3SAT problem. *Theor. Comput. Sci.*, 806:70–80, 2020.
- [12] D. Scheder and J. P. Steinberger. PPSZ for General k-SAT making Hertli's analysis simpler and 3-SAT faster. In R. O'Donnell, editor, 32nd Computational Complexity Conference, CCC 2017, July 6-9, 2017, Riga, Latvia, volume 79 of LIPIcs, pages 9:1–9:15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.
- [13] D. Scheder and N. Talebanfard. Super strong ETH is true for PPSZ with small resolution width. In S. Saraf, editor, 35th Computational Complexity Conference, CCC 2020, July 28-31, 2020, Saarbrücken, Germany (Virtual Conference), volume 169 of LIPIcs, pages 3:1–3:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [14] D. Scheder, B. Tang, S. Chen, and N. Talebanfard. Exponential lower bounds for the PPSZ k-sat algorithm. In S. Khanna, editor, Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013, pages 1253–1263. SIAM, 2013.
- [15] U. Schöning. A probabilistic algorithm for k-SAT and constraint satisfaction problems. In Proceedings of the 40th Annual Symposium on Foundations of Computer Science, pages 410–414. IEEE Computer Society, Los Alamitos, CA, 1999.

[16] N. Vyas and R. R. Williams. On super strong ETH. In M. Janota and I. Lynce, editors, Theory and Applications of Satisfiability Testing - SAT 2019 - 22nd International Conference, SAT 2019, Lisbon, Portugal, July 9-12, 2019, Proceedings, volume 11628 of Lecture Notes in Computer Science, pages 406–423. Springer, 2019.

A Proofs from Section 3, 4, and 6

Lemma A.1 (CCT similarity lemma, Lemma 33, restated). Let T_x be the canonical critical clause tree for T_x , u a node in T_x , v a descendant of u in T_x , and a := varlabel(u), b := varlabel(v). If v is canonical in T_x , then there is a corresponding node v' in T_a , the canonical clause tree of variable a, and the path from u to v has the same variable label sequence as the path from $\text{root}(T_a)$ to v' in T_a . In particular, varlabel(v') = b. Furthermore, v' is canonical.

Proof. Let $u = u_0, \ldots, u_t = v$ be the path from u to v in T_x , and $l_i = \text{varlabel}(u_i)$. Since v is canonical in T_x , each variable l_i on that path has exactly one canonical clause C_{l_i} . Furthermore, again by definition of canonical-ness, the assignment α_{u_i} violates C_{l_i} ; recall that α_{u_i} is the assignment obtained from $\alpha = (1, \ldots, 1)$ by flipping all variable labels on the path from $\text{root}(T_x)$ to u_i , including both endpoints.

We will now construct a corresponding path $\operatorname{root}(T_a) = u'_0, u'_1, \ldots, u'_t$ of canonical nodes in T_a such that $\operatorname{varlabel}(u'_i) = \operatorname{varlabel}(u_i)$. We start with $u'_0 := \operatorname{root}(T_a)$. Note that $\operatorname{varlabel}(u'_0) = \operatorname{varlabel}(u_0) = a$. Suppose u'_0, \ldots, u'_i have been defined, $i \leq l$. What is the clause label of u'_i ? Recall the assignment α_{u_i} , which violates C_{l_i} . What is $\alpha_{u'_i}$? It is the assignment obtained from $\alpha = (1, \ldots, 1)$ by flipping all variable labels on the path from $\operatorname{root}(T_a) = u'_0$ to u'_i . Thus, the set of flipped variables in $\alpha_{u'_i}$ is a subset of that in α_{u_i} ; in other words, $\alpha_{u'_i}(l) \geq \alpha_{u_i}(l)$ for every variable l. This means that $\alpha_{u'_i}$ violates C_{l_i} , as well. We conclude that the clause label of u'_i is C_{l_i} , and therefore u'_i is canonical, too.

If i = t, we are done; otherwise, we have to find a suitable u'_{i+1} in T_a . Write $C_{l_i} = (l_i \lor \bar{y} \lor \bar{z})$. The process constructing the critical clause trees creates two children for u_i and labels them y and z; without loss of generality, the child u_{i+1} has label y. Similarly, when constructing T_a , it uses C_l to create two children for u'_{i+1} and labels them y and z. We now set u'_{i+1} to be that child that has label y.

Lemma A.2 (Lemma 19, restated). There is some $c_{\text{PRIVILEGED}} > 0$, depending only on k, such that $\Pr[\text{Forced}(x,\pi)] \ge s_k + c_{\text{PRIVILEGED}} - o(1)$ for all privileged variables x, where o(1) converges to 0 as w grows.

Proof. We show that there are constants $c_1, c_2, c_3 > 0$, depending only on k, such that $\Pr[\operatorname{Forced}(x, \pi)] \geq s_k + c_i - o(1)$ whenever x is privileged due to reason (i) in Definition 18.

The case of privileged variables of type (1), i.e., those having at least two critical clauses, has already been addressed in the full versions of [3]. However, for the sake of completeness we will also discuss this case. We will introduce some operations on labeled trees T that never increase $\Pr[\operatorname{Cut}_r(T)]$. As a most simple example, suppose u is a node in T and not a safe leaf; form T' by adding a new child v to u. Then $\Pr[\operatorname{Cut}_r(T)] \geq \Pr[\operatorname{Cut}_r(T')]$, regardless of the label of v. This follows immediately from the definition of Cut_r .



Attaching additional descendants to u will not increase $\Pr[\operatorname{Cut}_r(T)]$.

This operation allows us to reduce case (3) to case (2). Indeed, suppose T_x has fewer than $(k-1)^2$ nodes at depth 2. Let Y_1, \ldots, Y_{k-1} be the children of the root of T_x and let y_1, \ldots, y_{k-1} be their respective labels. By assumption, some child Y_i has at most k-2children. Create a new node Z, attach it as an additional child to Y_i , and give it label y_1 . The resulting tree T' is a labeled tree, every node has at most k-1 children, and label y_1 occurs at depths 1 (at Y_1) and at depth 2 (at Z), so T' is of type (2).

Next, we (almost) reduce case (1) to case (2). Suppose x has two critical clauses, $C = (x \vee \bar{y}_1, \ldots, \bar{y}_{k-1})$ and $D = (x \vee \bar{z}_1, \ldots, \bar{z}_{k-1})$. Note that $k \leq |\{y_1, \ldots, y_{k-1}, z_1, \ldots, z_{k-1}\}| \leq 2(k-1)$. Suppose for the moment that it is less than 2(k-1), i.e., some variable y_i also appears in D. Without loss of generality, $y_1 = z_1$. Also, the two clauses are distinct, so let us assume that y_{k-1} does not appear in D. We will construct a (non-canonical) critical clause tree T'_x for x that is of type (2). Use C as clause label for the root. Note that this creates k-1 nodes Y_1, \ldots, Y_{k-1} at depth 1 with labels y_1, \ldots, y_{k-1} . The assignment label of Y_{k-1} is $\alpha[y_{k-1} \mapsto 0]$, which violates D; here we use the fact that D does not contain y_{k-1} . Thus, we can use D as clause label of node Y_{k-1} , which in turn creates k-1 nodes at depth 2, one of which has label z_1 . Now recall that $y_1 = z_1$ by assumption, so this label occurs once at depth 1 and somewhere at depth 2, and T_x is of type (2).

Summarizing, we are left with privileged variables of type (2) and those with two critical clauses C, D that share no variable besides x. Let us deal with type (2) first. We start with a proposition stating that assigning "fresh labels" to a node of T cannot increase $\Pr[\operatorname{Cut}_r(T)]$. This can be seen as an alternative proof of Lemma 7 in [8] that bypasses the FKG inequality for monotone Boolean functions (in fact, implicitly reproves it).

Proposition 76. Let T be a labeled tree and suppose the values $\{\pi(l)\}_{l \in L}$ are independent. Let u be a node in T with label z. Form a new tree T' by assigning u a fresh label z' and making $\pi(z')$ follow the same distribution as $\pi(z)$, but independent of everything else. Then $\Pr[\operatorname{Cut}(T)] \ge \Pr[\operatorname{Cut}(T')]$

Proof. We show that $\Pr[\operatorname{Cut}_r(T)] \geq \Pr[\operatorname{Cut}_r(T')]$ for all $r \in [0, 1]$. Let $\tau : L \setminus \{z, z'\} \rightarrow [0, 1]$ be a placement of all labels except z and z'. In fact, we claim that $\Pr[\operatorname{Cut}_r(T) \mid \tau] \geq \Pr[\operatorname{Cut}_r(T') \mid \tau]$ holds for all τ . Under the partial placement τ , the event $\operatorname{Cut}_r(T')$ becomes some monotone Boolean function f(b, b') in the Boolean variables $b := [\pi(z) < r]$ and $b' := [\pi(z) < r']$, and $\operatorname{Cut}_r(T)$ becomes f(b, b). This holds since T can be obtained from T' by merging the labels z and z'. Now under our distribution on placements, $\pi(z)$ and $\pi(z')$ follow the same distribution and therefore $\Pr[b = 1] = \Pr[b' = 1]$. This

means that $\Pr[f(b,b') = 1] = \Pr[f(b,b) = 1]$ unless f depends on both variables; the only monotone functions depending on both b and b' are $b \wedge b'$ and $b \vee b'$. If $f(b,b') = b \wedge b'$ then $\Pr[\operatorname{Cut}_r(T') \mid \tau] = \Pr[b \wedge b'] \leq \Pr[b] = \Pr[\operatorname{Cut}_r(T)]$ and our claim holds. Finally, $f(b,b') = b \vee b'$ cannot hold: the set of nodes in T' with label z or z' form an antichain, by Point 2 of Definition 9.

Now let T be a labeled tree to which (2) applies, i.e., some variable y appears at depth 1 and 2 in T_x . Let Y_1, Y_2 be those two nodes with label y. We apply the proposition to all nodes except Y_1 and Y_2 . Second, we add new children to nodes of depth less than h until all such nodes have exactly k-1 children, and all the $(k-1)^h$ nodes at depth h are safe leaves. Call this tree T. From the proposition and the discussion above, it follows that $\Pr[\operatorname{Cut}_r(T_x)] \geq \Pr[\operatorname{Cut}_r(T)]$. In a last step, give a fresh label y_1 to Y_1 and y_2 to Y_2 , and call the resulting tree T'. In T', all nodes have distinct labels, and we know what $\Pr[\operatorname{Cut}_r(T')]$ is: it is $Q_r^{(k-1)} - o(1)$.¹⁰ We also know that $\Pr[\operatorname{Cut}_r(T)] \geq \Pr[\operatorname{Cut}_r(T')]$ by the above proposition. Now, however, we have to take a closer look at the proof of the proposition since we want to show that $\Pr[\operatorname{Cut}_r(T)]$ is significantly larger than $\Pr[\operatorname{Cut}_r(T')]$.



The tree T': all labels are distinct.

In T', we denote the node with label w_2 by W_2 ; that with label z_2 by Z_2 , and so on. Since T' and T has the same node set, we use this notation for the nodes in T, too. Furthermore, for a node u, we denote the subtree of T (or T') rooted at node u by T_u (or T'_u). As in the proof of the proposition, we fix some partial placement $\tau : L \setminus \{y_1, y_2, y\} \rightarrow [0, 1]$. As we have seen in the proof, $\Pr[\operatorname{Cut}_r(T) \mid \tau] \geq \Pr[\operatorname{Cut}_r(T') \mid \tau]$ holds for every such τ . Call τ good if the following holds: (1) $\pi(w_2) \geq r$; (2) $\pi(z_2), \ldots, \pi(z_{k-1}), \pi(w_3), \ldots, \pi(w_{k-1}) < r$; (3) $\neg \operatorname{Cut}_r(T_{Y_1})$ and $\neg \operatorname{Cut}_r(T_{Y_2})$. The events described in (1–3) are independent; those in (1) and (2) happen with probability exactly $(1 - r)r^{2k-5}$. Those in (3) happen with probability at least $(1 - Q_r)^2$.

Under a good τ , the $\operatorname{Cut}_r(T')$ becomes $[\pi(y_1) < r \land \pi(y_2) < r]$ and has probability r^2 , and $\operatorname{Cut}_r(T)$ becomes $[\pi(y) < r]$, which has probability r. Therefore,

$$\begin{aligned} \Pr[\operatorname{Cut}_r(T)] - \Pr[\operatorname{Cut}_r(T')] &\geq \Pr[\tau \text{ is good}] \cdot (r - r^2) \\ &\geq (1 - r)r^{2k - 5}(1 - Q_r)^2(r - r^2) = (1 - r)^2 r^{2k - 4}(1 - Q_r)^2 . \end{aligned}$$

Putting everything together and integrating over r, we conclude that

$$\Pr[\operatorname{Cut}(T_x)] \ge s_k - o(1) + \int_0^1 (1-r)^2 r^{2k-4} (1-Q_r)^2 \, dr \; .$$

¹⁰We will write Q_r instead of $Q_r^{(k-1)}$ from now on since k is understood.

It is clear that the integral is some positive constant depending solely on k.

We are left with the case that x has two critical clauses $C = (x \vee \bar{y}_1, \ldots, \bar{y}_{k-1})$ and $D = (x \vee \bar{z}_1, \ldots, \bar{z}_{k-1})$, and $y_i \neq z_j$ for all $1 \leq i, j \leq k-1$. It is clear that x is forced if all y_i come before x or all z_i come before x. Therefore,

$$\Pr[\operatorname{Forced}(x,\pi) \mid \pi(x) = r] \ge r^{k-1} + r^{k-1} - \Pr[\operatorname{all} y_i \text{ and all } z_i \text{ come before } x]$$
$$\ge 2 r^{k-1} - r^{2k-2}.$$

On the other hand, by focusing solely on the canonical critical clause tree of x, we can apply Lemma 13 and conclude that

$$\Pr[\operatorname{Forced}(x,\pi) \mid \pi(x) = r] \ge Q_r - o(1) ,$$

(we write Q_r instead of $Q_r^{(k)}$ since k is understood), and therefore

$$\Pr[\operatorname{Forced}(x,\pi) = 1] \ge \int_0^1 \max(2r^{k-1} - r^{2k-2}, Q_r) \, dr - o(1)$$
$$= s_k - o(1) + \int_0^1 \max\left(0, 2r^{k-1} - r^{2k-2} - Q_r\right) \, dr$$

It remains to show that the latter term is positive for a substantial range of $r \in [0, 1]$. We claim that if r is sufficiently small, Q_r is only a tiny factor larger than r^{k-1} . Indeed, From Proposition 12, we know that $Q_r \leq e r^{k-1}$, thus $P_r = r \vee Q_r = r + (1-r)e r^{k-1} = r(1+e(1-r)r^{k-2})$ and in turn $Q_r = P_r^{k-1} \leq (r+e r^{k-1})^{k-1} = r^{k-1}(1+e r^{k-2})^{k-1} < r^{k-1}e^{e(k-1)r^{k-2}}$. We check that $e^{e(k-1)r^{k-2}} \leq 1.5$ for all $k \geq 3$ and $r \leq 1/16$. Therefore,

$$\int_{0}^{1} \max\left(0, 2r^{k-1} - r^{2k-2} - Q_r\right) dr \le \int_{0}^{1/16} \left(2r^{k-1} - r^{2k-2} - 1.5r^{k-1}\right) dr$$
$$= \int_{0}^{1/16} \left(\frac{1}{2}r^{k-1} - r^{2k-2}\right) dr$$

and the latter is some positive constant c_1 depending only on k.

Lemma A.3 (Lemma 34, restated). There is a set $H \subseteq E(SG)$ of maximum degree 2 (i.e., H consists of paths and cycles) with $|H| \ge n - |ID_1| - 2 |ID_0|$.

Proof of Lemma 34. For each vertex x in SG with $\deg_{SG}(x) \ge 3$, mark $\deg_{SG}(x) - 2$ of its incident edges (choose them arbitrarily). Let H be the set of unmarked edges. The number of marked edges is at most

$$\sum_{x:\deg_{SG}(x)\geq 3} (\deg_{SG}(x) - 2)$$

The total number of edges is n, thus $\sum_{x \in V} \deg_{SG}(x) = n$, and

$$0 = \sum_{x} (\deg(x) - 2) = \sum_{x: \deg_{SG}(x) \ge 3} (\deg_{SG}(x) - 2) - |\mathrm{ID}_1| - 2 |\mathrm{ID}_0| ,$$

We conclude that the number marked edges is at most

$$\sum_{\text{:deg}_{SG}(x) \ge 3} (\text{deg}_{SG}(x) - 2) = |\text{ID}_1| + 2 |\text{ID}_0| .$$

Thus, $|H| \ge n - |ID_1| - 2 |ID_0|$, and the maximum degree in H is obviously at most 2. \Box

B Proofs concerning the construction of D

Theorem B.1 (Theorem 40, restated). Let $A_v \subseteq [0,1]$ for $v \in V(G)$ be non-empty intervals; let $V(G) = K \uplus I$ and $E_I := \{\{u, v\} \in E \mid u, v \in I\}$. Then

$$\Pr[X_k \in A_k \ \forall k \in K \mid X_i \in A_i \ \forall i \in I] = \prod_{k \in K} \mu(A_k) \cdot \left(1 + \epsilon \frac{\sum_{\{u,v\} \in E \setminus E_I} T_u T_v}{1 + \epsilon \sum_{\{u,v\} \in E_I} T_u T_v}\right) ,$$

where μ is the Lebesgue measure on [0,1] and $T_u := \mathbb{E}_{x \in A_u}[\phi(x)]$, the expectation being taken with respect to the uniform distribution on A_u .

Proof. If some $A_k, k \in K$ has zero measure $(A_k \text{ is a point})$ then both sides are 0. For simplicity, let's assume all intervals have positive measure, including the $A_i, i \in I$; the case $A_i = \{r\}$ can then be obtained by taking a limit.

As a small notational trick, define A'_v to be [0,1] if $v \in K$ and A_v if $v \in I$. Then the conditional probability is the fraction

$$\frac{\Pr[X_v \in A_v \ \forall v \in V]}{\Pr[X_v \in A'_v \ \forall v \in V]}$$

Let us work on the numerator. We define the Cartesian product $A := \prod_{v \in V} A_v$ and let μ be the Lebesgue measure on [0, 1] and $[0, 1]^d$ in general. Note that $T_u = \mathbb{E}_{x \in A_u}[\phi(x)] = \int_{A_u} \phi(x) dx / \mu(A_u)$. Then

$$\frac{\Pr[X_v \in A_v \ \forall v \in V]}{\mu(A)} = \frac{\int_A (1 + \epsilon \sum_{\{u,v\} \in E} \phi(x_u)\phi(x_v) \, d\mathbf{x})}{\mu(A)}$$
$$= \underset{\mathbf{x} \in A}{\mathbb{E}} \left(1 + \epsilon \sum_{\{u,v\} \in E} \phi(x_u)\phi(x_v) \right)$$
$$= 1 + \epsilon \sum_{\{u,v\} \in E} T_u T_v \, .$$

Let us work on the denominator similarly. We define $A' := \prod_{v \in V} A'_v$ and $T'_u := \mathbb{E}_{x \in A'_u}[\phi(u)]$. Note that $T'_u = T_u$ if $u \in I$ and $T'_u = 0$ if $u \in K$. We get

$$\frac{\Pr[X_v \in A'_v \ \forall v \in V]}{\mu(A')} = 1 + \epsilon \sum_{\{u,v\} \in E} T'_u T'_v$$

Finally, the conditional probability is

$$\frac{\Pr[X_v \in A_v \ \forall v \in V]}{\Pr[X_v \in A'_v \ \forall v \in V]} = \frac{1 + \epsilon \sum_{\{u,v\} \in E} T_u T_v}{1 + \epsilon \sum_{\{u,v\} \in E} T'_u T'_v} \cdot \frac{\mu(A)}{\mu(A')}$$

$$= \prod_{k \in K} \mu(A_k) \cdot \frac{1 + \epsilon \sum_{\{u,v\} \in E} T_u T_v}{1 + \epsilon \sum_{\{u,v\} \in E} T'_u T'_v}$$

$$= \prod_{k \in K} \mu(A_k) \cdot \frac{1 + \epsilon \sum_{\{u,v\} \in E} T_u T_v}{1 + \epsilon \sum_{\{u,v\} \in E_I} T_u T_v} \quad (\text{since } T'_u = T_u \cdot \mathbf{1}_{u \in I})$$

$$= \prod_{k \in K} \mu(A_k) \cdot \left(1 + \frac{\epsilon \sum_{\{u,v\} \in E \setminus E_I} T_u T_v}{1 + \epsilon \sum_{\{u,v\} \in E_I} T_u T_v}\right)$$

This completes the proof of the theorem.

Corollary B.2 (Corollary 42, restated). Let $u \in V(G)$ and $A_u := [0, r]$; for each other $v \in V \setminus \{u\}$, suppose A_v is one of $\{r\}$, [0, r], [r, 1], or [0, 1]. Define $T_v^- = \min(0, T_v)$. Then

$$\Pr[X_u \in A_u \mid X_v \in A_v \text{ for all other } v] \ge r + \frac{\epsilon \gamma(r) \sum_{v:\{u,v\} \in E} T_v^-}{1 - \frac{2}{25} \epsilon(|E| - 1)} .$$
(38)

holds for our particular choice $\gamma(r) = r(1-2r)^{3/2}$. Furthermore, if $|E(G)| \leq 17$ and $\epsilon \leq 0.13$, this is at least

$$r+1.2\,\epsilon\gamma(r)\sum_{v:\{u,v\}\in E}T_v^-\ .$$

Proof. We apply Theorem 40 with $K = \{u\}$ and $I = V \setminus \{u\}$ and A_v as in the corollary and see that

$$\Pr[X_u \in A_u \mid X_v \in A_v \; \forall v \in I] = r \left(1 + \epsilon \frac{\sum_{v:\{u,v\} \in E} T_u T_v}{1 + \epsilon \sum_{\{v,w\} \in E_I} T_v T_w} \right)$$
$$= r + \frac{\epsilon \gamma(r) \sum_{v:\{u,v\} \in E} T_v}{1 + \epsilon \sum_{\{v,w\} \in E_I} T_v T_w} \quad (\text{since } T_u = \frac{\gamma(r)}{r})$$
$$\geq r + \frac{\epsilon \gamma(r) \sum_{v:\{u,v\} \in E} T_v^-}{1 + \epsilon \sum_{\{v,w\} \in E_I} T_v T_w} .$$

If u is an isolated vertex in G than this is obviously equal to r and the theorem holds. Otherwise, observe that the the numerator is at most 0; it remains to bound the denominator from below. Consider a single term $T_v T_w$. There are four possible choices for A_v , and similar for A_w . Note that T_v is $\phi(r)$, $\frac{\gamma(r)}{r}$, $\frac{-\gamma(r)}{1-r}$, and 0 if A_v is $\{r\}$, [0, r], [r, 1], and [0, 1], respectively. This makes a total of sixteen combinations; however, most of them are obviously non-negative (for example if $A_v = A_w$ or one of them is [0, 1]); there are only three values for $T_v T_w$ that are not obviously non-negative:

- 1. $T_v T_w = \frac{\gamma(r)}{r} \cdot \frac{-\gamma(r)}{1-r}$, if $A_v = [0, r]$ and $A_w = [r, 1]$, or vice versa;
- 2. $T_v T_w = \phi(r) \cdot \frac{\gamma(r)}{r}$, if $A_v = \{r\}$ and $A_w = [0, r]$, or vice versa;
- 3. $T_v T_w = \phi(r) \cdot \frac{-\gamma(r)}{1-r}$, if $A_v = \{r\}$ and $A_w = [r, 1]$, or vice versa.

One checks that all three functions are at least -2/25, with equality achieved by $\phi(r) \cdot \frac{\gamma(r)}{r}$ and r = 3/10. There are at most |E| - 1 terms in the denominator since u is incident to at least one $v \in I$, and therefore the denominator is at least $1 - \frac{2}{25} \epsilon(|E| - 1)$. This completes the proof of Corollary 42.

Lemma B.3 (Lemma 43, restated). If $|\phi(x)| \leq 1$ for all $x \in [0,1]$ then $\operatorname{KL}(D_{\epsilon}^{\gamma}||U) \leq \frac{m_2}{\ln(2)} \cdot f_{\mathrm{KL}}(\epsilon)$. Using $\ln(1-\epsilon) \leq -\epsilon - \epsilon^2/2$, this is at most to $\frac{m_2}{2\ln(2)}(\epsilon^2 + \epsilon^3)$.

Proof. By definition, if $f_Q(z), f_P(z)$ are the probability densities of Q and P, then

$$\mathbb{D}(P||Q) = \mathop{\mathbb{E}}_{z \sim P} \left[\log_2 \left(\frac{f_P(z)}{f_Q(z)} \right) \right] \; .$$

In our case, Q is the uniform distribution on [0, 1] and P has density $f_P(x) = 1 + \epsilon \phi(x)$ The divergence is

$$\mathbb{D}\left(D_{\epsilon}^{\gamma}||U\right) = \frac{1}{\ln(2)} \int_{x} (1 + \epsilon\phi(x)) \ln(1 + \epsilon\phi(x)) \, dx \tag{39}$$

Write $t := \epsilon \phi(x)$. Using the Taylor expansion of $\ln(1+t)$, we see that

$$(1+t)\ln(1+t) = (1+t)\sum_{n=1}^{\infty} (-1)^{n+1} \frac{t^n}{n}$$
$$= \sum_{n=1}^{\infty} (-1)^{n+1} \frac{t^n}{n} + \sum_{n=1}^{\infty} (-1)^{n+1} \frac{t^{n+1}}{n}$$
$$= \sum_{n=1}^{\infty} (-1)^{n+1} \frac{t^n}{n} + \sum_{k=2}^{\infty} (-1)^k \frac{t^k}{k-1}$$
$$= t + \sum_{n=2}^{\infty} t^n \cdot (-1)^n \cdot \left(\frac{1}{n-1} - \frac{1}{n}\right)$$
$$= t + \sum_{n=2}^{\infty} \frac{(-1)^n}{n(n-1)} \cdot t^n .$$

Thus, (39) becomes

$$\frac{1}{\ln(2)} \int_{x} \epsilon \phi(x) dx + \sum_{n=2}^{\infty} \frac{(-1)^{n}}{\ln(2) n(n-1)} \int_{x} \epsilon^{n} \phi^{n}(x) dx$$

The first sum is $\epsilon m_1 = 0$. The second sum equals

$$\begin{split} &\sum_{n=2}^{\infty} \frac{(-1)^n \epsilon^n}{\ln(2) n(n-1)} \int_x \phi^n(x) \, dx \\ &\leq \sum_{n=2}^{\infty} \frac{\epsilon^n}{\ln(2) n(n-1)} \int_x |\phi^n(x)| \, dx \\ &\leq \sum_{n=2}^{\infty} \frac{\epsilon^n}{\ln(2) n(n-1)} \int_x |\phi^2(x)| \, dx \qquad (\text{since } \phi(x) \leq 1) \\ &= \frac{m_2}{\ln(2)} \sum_{n=2}^{\infty} \frac{\epsilon^n}{n(n-1)} \, . \end{split}$$

It remains to find a closed form for $\sum_{n=2}^{\infty} \frac{x^n}{n(n-1)}$.

Claim.
$$\sum_{n=2}^{\infty} \frac{x^n}{n(n-1)} = (1-x)\ln(1-x) + x = f_{\text{KL}}(x)$$
 for all $|x| < 1$.

Proof. Both sides converge / are defined for |x| < 1. We compare the Taylor series. For x = 0, both sides vanish, so the constant terms agree. Differentiating both sides, we get $\sum_{n=2}^{\infty} \frac{x^{n-1}}{n-1}$ on the left side and $-\ln(1-x)$ on the right side. Both first derivatives vanish for x = 0, so the linear terms agree as well. Finally, differentiating once more, the right side becomes $\frac{1}{1-x}$ and the left side becomes $\sum_{n=2}^{\infty} x^{n-2}$. This is a geometric series and equals $\frac{1}{1-x}$ for all |x| < 1.

Lemma B.4 (Lemma 44, restated). Let G be a cycle or a path, consisting of at most t edges. For $\gamma(r) = r(1-2r)^{3/2}$, $\phi(r) = \gamma'(r)$, $\epsilon \leq 0.13$, and $t \leq 17$, it holds that $\operatorname{KL}(D^G||U) \leq 0.0064 \epsilon^2 t$.

Proof. This is more challenging than if D were the Markov chain outlined above. We use the estimate $\ln(1+z) \leq z - z^2/2 + z^3/3$, which holds for all $z \in \mathbb{R}$, and write $z = z(\mathbf{x}) = \epsilon \sum_{i=1}^{t} \phi(x_{i-1})\phi(x_i)$ for brevity. By definition of KL, we have

$$\ln(2) \text{KL}(D||U) = \int (1+z(\mathbf{x})) \ln(1+z(\mathbf{x})) d\mathbf{x}$$

$$\leq \int (1+z)(z-z^2/2+z^3/3) d\mathbf{x}$$

$$= \int \left(z+\frac{z^2}{2}-\frac{z^3}{6}+\frac{z^4}{3}\right) d\mathbf{x} ,$$

where the integral is always over the whole space $[0, 1]^V$. We will therefore replace \int by \mathbb{E} , referring to the distribution that samples each x_i independently and uniformly over [0, 1]. This change is of course purely notational but we feel that \mathbb{E} is slightly more intuitive than \int . We define $m_d := \mathbb{E}[\phi^d(x)]$ and check that $m_2 = \frac{3}{32}$, $m_3 = \frac{12}{385}$, and $m_4 = \frac{9}{224}$, for our choice of $\gamma(r) = r(1-2r)^{3/2}$. We integrate / take the expectation of each of the four terms above separately. For the linear term z we get

$$\mathbb{E}[z(\mathbf{x})] = \mathbb{E}\left[\epsilon \sum_{\{u,v\}\in E} \phi(x_u)\phi(x_v)\right] = \epsilon \sum_{\{u,v\}\in E} \mathbb{E}[\phi(x_u)] \mathbb{E}[\phi(x_v)] = 0.$$

For $z^2/2$ we get

$$\mathbb{E}\left[\frac{z^2}{2}\right] = \frac{\epsilon^2}{2} \mathbb{E}\left[\left(\sum_{\{u,v\}\in E} \phi(x_u)\phi(x_v)\right)^2\right]$$
$$= \frac{\epsilon^2}{2} \sum_{\{u,v\},\{u',v'\}\in E} \mathbb{E}\left[\phi(x_u)\phi(x_v)\phi(x_{u'})\phi(x_{v'})\right] .$$

If $\{u, v\} \neq \{u', v'\}$ then $u \notin \{u', v'\}$ without loss of generality, and $\phi(x_u)$ is independent of the other three terms in the expectation; thus, we can factor $\mathbb{E}[\phi(x_u)]$ out and see that the whole term vanishes. Thus, only terms for $\{u, v\} = \{u', v'\}$ remain:

$$\frac{\epsilon^2}{2} \sum_{\{u,v\}\in E} \mathbb{E} \left[\phi(x_u)^2 \phi(x_v)^2\right]$$
$$\frac{\epsilon^2}{2} \sum_{\{u,v\}\in E} \mathbb{E} \left[\phi(x_u)^2\right] \mathbb{E} \left[\phi(x_v)^2\right]$$
$$= \frac{\epsilon^2}{2} t \left(\mathbb{E} \left[\phi(x)^2\right]\right)^2 = \frac{\epsilon^2 t m_2^2}{2} .$$

To facilitate computation of the cubic and quartic terms below, let's put this argument in a more general context: for a tuple $F = (e_1, \ldots, e_i)$ of edges in E, let $T_F :=$

 $\mathbb{E}\left[\prod_{i=1}^{i}\prod_{u\in e_{i}}\phi(x_{u})\right]$. Let deg_F(u) denote the number of edges e in F incident to u. We calculate

$$T_F = \mathbb{E}\left[\prod_{e \in F} \prod_{u \in e} \phi(x_i)\right]$$
$$= \mathbb{E}\left[\prod_{u \in V(G)} (\phi(x_u))^{\deg_F(u)}\right]$$
$$= \prod_{u \in V(G)} \mathbb{E}\left[(\phi(x_u))^{\deg_F(u)}\right]$$
$$= \prod_{u \in V(G)} m_{\deg_F(u)} .$$

Note that $m_1 = \mathbb{E}[\phi(x)] = 0$; let us say vertex *i* is exposed in *F* if $\deg_F(i) = 1$, thus $T_F = 0$ if *F* has an exposed vertex. In the above calculation of $\mathbb{E}\left[\frac{z^2}{2}\right]$, the multiset *F* consisted of only two edges, *e* and *e'*. It is obvious that *F* has an exposed vertex unless e = e'.

Next, let us analyze the degree-3-terms $(-z^3/6)$:

$$\mathbb{E}\left[\frac{-z^3}{6}\right] = \frac{-\epsilon^3}{6} \mathbb{E}\left[\left(\sum_{\{u,v\}\in E}^t \phi(x_u)\phi(x_v)\right)^3\right] = \frac{-\epsilon^3}{6} \sum_{e,f,g\in E} T_{(e,f,g)}$$

What is $T_{(e,f,g)}$? If $e = f = g = \{u, v\}$ then u and v have degree 3 in F, and $T_F = m_3^2$. If G is a triangle and e, f, g are all distinct (for which there are 3! = 6 possibilities), then all $u \in V$ have deg_F(u) = 2, and $T_F = m_2^3$. Otherwise, F has an exposed vertex and $T_F = 0$. We conclude that

$$\mathbb{E}\left[\frac{-z^3}{6}\right] = \begin{cases} \frac{-\epsilon^3}{6} \left(3 m_3^2 + 6 m_2^3\right) & \text{if } G \text{ is a triangle,} \\ \frac{-\epsilon^3 t m_3^2}{6} & \text{else.} \end{cases}$$

In both cases, this is negative. Finally, $z^4/3$:

$$E\left[\frac{z^4}{3}\right] = \frac{\epsilon^4}{3} \sum_{e,f,g,h} T_{(e,f,g,h)}$$

Case 1. G is not a triangle and not a 4-cycle.

It is easy to see that F = (e, f, g, h) is either of the form (1) (e, e, e, e), i.e., contains one edge four times or (2) (e, f, e, f) for $e \neq f$ or a permutation thereof, or (3) has an exposed vertex. Form (1) occurs t times (once for every e) and contributes a total of $t m_4^2$, since u and v have degree 4 in F, for $e = \{u, v\}$. For each set $\{e, f\}$ of two edges, form (2) appears six times (choose the two positions of e). This makes a total of $\binom{t}{2} \cdot 6 = 3t(t-1)$ times form (2) appears. Form (2) can, however, appear in two sub-forms: either (2.1) the edges e, f can be incident or (2.2) they are not. Note that the number of sets $\{e, f\}$ that are of form (2.1) is t - 1 if G is a path and t if G is a cycle. If term $T_{e,f,g,h}$ of form (2.1), say $T_{e,e,f,f}$ with $e = \{u, v\}$ and $f = \{v, w\}$ then $\deg_F(u) = \deg_F(w) = 2$ and $\deg_F(v) = 4$ and therefore $T_{e,e,f,f} = m_2^2 m_4$. If it is of form (2.2) then four vertices are involved, each of degree 2, and $T_{e,e,f,f} = m_2^4$. Therefore,

$$\mathbb{E}\left[\frac{z^4}{3}\right] = \frac{\epsilon^4}{3} \cdot \begin{cases} tm_4^2 + 3(t-1)\left(2\,m_4m_2^2 + (t-2)m_2^4\right) & \text{if } G \text{ is a } t\text{-path}\\ tm_4^2 + 3\,t\left(2\,m_4m_2^2 + (t-3)m_2^4\right) & \text{if } G \text{ is a } t\text{-cycle} \end{cases}$$

One checks that the expression for the *t*-cycle is at least the expression for the *t*-path; the key inequality here is $m_2^2 \leq m_4$, which follows from setting $Y := \phi^2(x)$ and observing that $m_2^2 = (\mathbb{E}[Y])^2 \leq \mathbb{E}[Y^2] = m_4$. We conclude that

$$\mathbb{E}\left[\frac{z^4}{3}\right] \le \frac{\epsilon^4 t}{3} \cdot \left(m_4^2 + 6 \, m_4 m_2^2 + 3 \, (t-3) m_2^4\right) \ . \tag{40}$$

Adding the quadratic and quartic term (the cubic one is negative), and using the fact that $t \leq 17$, we conclude that

$$\ln(2)\mathrm{KL}(D||U) \leq \frac{\epsilon^2 t m_2^2}{2} + \frac{\epsilon^4 t}{3} \cdot \left(m_4^2 + 6 \, m_4 m_2^2 + 42 \, m_2^4\right) \qquad (\text{since } t \leq 17)$$
$$= \epsilon^2 t \left(\frac{m_2^2}{2} + \epsilon^2 \, \frac{m_4^2 + 6 \, m_4 m_2^2 + 42 \, m_2^4}{3}\right) \qquad (41)$$
$$\leq 0.004434 \, \epsilon^2 t \;,$$

where the last inequality uses $\epsilon \leq 0.13$. This implies that $\text{KL}(D||U) \leq 0.0064 \epsilon^2 t$.

Case 2. G is a triangle. Then we get the same as in (40) plus some cases not accounted for. If (3) F := (e, f, g, h) = (e, e, e, f) or a permutation therefore, then one endpoint of f is exposed and $T_F = 0$. If (4) F := (e, f, g, h) = (e, e, f, g) then no endpoint is exposed (since e, f, g are the three edges of the triangle). The three points have degree 3, 3, and 2 in F, respectively, and thus $T_F = m_3^2 m_2$. How many such tuples F are there? There are three ways to choose one edge to occur twice in F. Once this has been chosen, say e appears twice, it remains to choose where in the 4-tuple F the two other edges, f and g appear; there are $4 \cdot 3 = 12$ choices for this, making a total of 36 tuples F in which exactly one edge of e, f, g appears twice. Thus,

$$\mathbb{E}\left[\frac{z^4}{3}\right] = \frac{\epsilon^4}{3} \left(3 \, m_4^2 + 18 m_4 m_2^2 + 36 \, m_3^2 m_2\right)$$
$$= \frac{\epsilon^4 t}{3} (m_4^2 + 6 \, m_4 m_2^2 + 12 \, m_3^2 m_2) . \qquad \text{(if } G \text{ is triangle)}$$

Adding things up, we see that

$$\ln(2) \text{KL}(D||U) \leq \frac{\epsilon^2 t m_2^2}{2} + \frac{\epsilon^4 t}{3} (m_4^2 + 6 \, m_4 m_2^2 + 12 \, m_3^2 m_2) = \epsilon^2 t \left(\frac{m_2^2}{2} + \epsilon^2 \, \frac{m_4^2 + 6 \, m_4 m_2^2 + 12 \, m_3^2 m_2}{3} \right) \leq 0.00443 \, \epsilon^2 t \;, \qquad (\text{since } \epsilon \leq 0.13)$$

which again implies $\operatorname{KL}(D||U) \leq 0.0064 \epsilon^2 t$.

Case 3. G is a four-cycle. In addition (40), we get some extra terms as well. If (3) F := (e, f, g, h) = (e, e, e, f) or some permutation thereof, i.e., some edge has multiplicity

3, then f has at least one exposed vertex, and $T_F = 0$. If (4) exactly one edge has multiplicity 2, so F = (e, e, f, g) or a permutation thereof then F has at least one exposed vertex, and $T_F = 0$, as well. The case F = (e, e, f, f), i.e., two edges appear twice each, has already been accounted for in (40) as form (2). We get an additional term if (5) all edges are distinct: F = (e, f, g, h); in this case, all vertices have degree 2, and $T_F = m_2^4$. There are of course 4! tuples of this form, and thus

$$\mathbb{E}\left[\frac{z^4}{3}\right] = \frac{\epsilon^4}{3} \left(4 \, m_4^2 + 24 \, m_4 m_2^2 + 12 \, m_2^4 + 24 \, m_2^4\right) \qquad \text{(if } G \text{ is 4-cycle)}$$

Altogether, if G is a 4-cycle we get that

$$\ln(2)\mathrm{KL}(D||U) \leq \frac{\epsilon^2 t m_2^2}{2} + \frac{\epsilon^4}{3} \left(4 m_4^2 + 24 m_4 m_2^2 + 36 m_2^4\right)$$
$$= \epsilon^2 t \left(\frac{m_2^2}{2} + \epsilon^2 \frac{m_4^2 + 6 m_4 m_2^2 + 9 m_2^4}{3}\right)$$
$$\leq \text{ the expression in (41)}.$$

Thus, the bound in the lemma also holds for the 4-cycle. This concludes the proof. \Box

C Proofs for Section 7

Lemma C.1 (Extending T to infinity; Lemma 48, restated). Let T be a labeled tree of height h in which all safe leaves have depth h, and for every node v of depth greater than h', the following hold: (1) its label l_v is not shared by any other node; (2) $\pi(l_v)$ is independent of everything else; (3) $r_{\min} \leq \Pr[\pi(l_v) < r] \leq r_{\max}$ (think of r_{\min} and r_{\max} being close to r); (4) v has exactly two children if $d(v) \leq h - 1$, and it is a safe leaf if d(v) = h.

Construct T' from T by replacing each safe leaf v by a copy of T_{∞} , so T' has no safe leaves anymore. All new nodes get $\delta_v = 0$, i.e., $\Pr[\pi(l) < r] = r$ for all new labels. Then

$$\Pr[\operatorname{Cut}_r(T)] \ge \Pr[\operatorname{Cut}_r(T')] - 2^h \tilde{r}^{h-h'}$$

Proof. Let $T_0 := T$. Enumerate the leaves of T_0 as $i = 1, \ldots, L$, with $L \leq 2^h$. Form T_{i+1} from T_i by replacing the i^{th} leaf with a copy of T_{∞} , and let $T' := T_L$.

Claim. For $r \in [0, 1/2]$, $\Pr_D[\operatorname{Cut}_r(T_i)] \ge \Pr_D[\operatorname{Cut}_r(T_{i+1})] - \tilde{r}^{h-h'}$.

Proof. Let v be the $(i + 1)^{\text{th}}$ leaf of T. That is, v is a leaf in T_i but the root of a copy of T_{∞} in T_{∞} . Besides that, T_i and T_{i+1} are identical. Let τ be an assignment to all labels of T_i except v. Note that under τ , the events $\operatorname{Cut}_r(T_i)$ and $\operatorname{Cut}_r(T_{i+1})$ either (1) both become \emptyset , i.e., neither happens regardless of $\pi(v)$; (2) both become Ω , i.e., both happen regardless of $\pi(v)$; or (3) $\operatorname{Cut}_r(T_i)$ becomes $[\pi(v) < r]$ and $\operatorname{Cut}_r(T_{i+1})$ becomes w $\operatorname{Cut}(T_v)$, where T_v is the copy of T_{∞} rooted at v in T_{i+1} . Say v is pivotal under τ if (3) happens. Observe that

$$\Pr_{D}[\operatorname{Cut}_{r}(T_{i+1})] - \Pr_{D}[\operatorname{Cut}_{r}(T_{i})] = \Pr_{\tau \sim D}[v \text{ is pivotal under } \tau] \cdot \left(\Pr_{D}[\operatorname{wCut}(T_{v})] - \Pr_{D}[\pi(v) < r]\right)$$
$$\leq \Pr_{\tau \sim D}[v \text{ is pivotal under } \tau] \cdot (P_{\tilde{r}} - \tilde{r}) \quad .$$
(42)

What has to happen for v to be pivotal under τ ? Let root = $u_0, u_1, \ldots, u_h = v$ be the path from the root to v; for $h' + 2 \le i \le h$, let w_i be the child of u_{i-1} that is not u_i , i.e.,

 w_i is an "aunt" of v. This exists since u_{i-1} has depth at least h' and thus has exactly two children. Let T_{w_i} be the subtree of T rooted at w_i . Observe that for v to be pivotal under τ , the following are necessary:

- 1. $\pi(u_{h'+1}), \ldots, \pi(u_{h-1}) \ge r$ and
- 2. wCut_r $(T_{w_{h'+2}}),\ldots,$ wCut_r (T_{w_h})

These events are independent and happen with probability

$$\Pr[v \text{ is pivotal under } \tau] = \prod_{i=h'+1}^{h-1} \Pr[\pi(u_i) \ge r] \cdot \prod_{i=h'+2}^{h} \Pr[\operatorname{wCut}_r(T_{w_i})]$$
$$= (1-\tilde{r})^{h-h'-1} \cdot \prod_{i=h'+2}^{h} \Pr[\operatorname{wCut}_r(T_{w_i})]$$
$$\le (1-\tilde{r})^{h-h'-1} \cdot (P_{\tilde{r}})^{h-h'-1}$$
$$= \tilde{r}^{h-h'-1}.$$

To justify the inequality, note that in T_{w_i} , every node either has two children or is a safe leaf. Thus, the event wCut_r is less likely in T_{w_i} than in the infinite binary tree, where it is $P_{\tilde{r}}$. The last inequality follows because $P_{\tilde{r}} = \frac{\tilde{r}}{1-\tilde{r}}$ whenever $\tilde{r} \leq 1/2$. We see that

$$(42) \le \tilde{r}^{h-h'-1}(P_{\tilde{r}} - \tilde{r}) = \tilde{r}^{h-h'} \cdot \frac{r}{1 - \tilde{r}} \le \tilde{r}^{h-h'}$$

This proves the claim.

Iterating the claim over all $L \leq 2^h$ leaves, we see that $\Pr_D[\operatorname{Cut}_r(T_0)] \geq \Pr_D[\operatorname{Cut}_r(T_L)] - 2^h \tilde{r}^{h-h'}$, which proves the lemma.

 \square

Lemma C.2 (Biased node lemma, Lemma 49, restated). If v is a maximal node in W, *i.e.*, no proper ancestor of v is in W, then

$$\Pr_{D}[\operatorname{Cut}_{r}(T)] \geq \Pr_{D_{v}}[\operatorname{Cut}_{r}(T)] - \delta_{v} \cdot \frac{1 - Q_{r-\delta_{\max}}}{1 - r} \cdot r^{d}$$

where $d := d_T(v)$.

Proof of the lemma. Conceptually, this is similar to the proof of the claim within the proof of Lemma 48. Let τ be a placement to all labels except v.¹¹ Note that the distribution of τ is identical under D and under D_v . Under the partial placement τ , the event $\operatorname{Cut}_r(T)$ either (1) happens regardless of $\pi(v)$, (2) does not happen, regardless of $\pi(v)$, or (3) becomes $[\pi(v) < r]$. We say v is pivotal under τ if (3) happens. If v is pivotal under τ then $\Pr_D[\operatorname{Cut}_r(T) \mid \tau] = r - \delta_v$ and $\Pr_{D_v}[\operatorname{Cut}_r(T) \mid \tau] = r$. Therefore,

$$\Pr_{D_v}[\operatorname{Cut}_r(T)] - \Pr_D[\operatorname{Cut}_r(T)] = \delta_v \Pr_{\tau \sim D}[v \text{ is pivotal under } \tau] .$$

We will now prove an upper bound on $\Pr_{\tau \sim D}[v \text{ is pivotal under } \tau]$. As before, let root = $u_0, u_1, \ldots, u_d = v$ be the path from the root to v. For $1 \leq i \leq d$, let w_i be the child of u_{i-1} that is not u_i , i.e., w_i is an aunt of v, and let T_{w_i} be the subtree of T rooted at w_i . We observe that v is pivotal under τ if and only if

 $^{^{11}{\}rm Strictly}$ speaking, v is a node, not a label; however, since all labels are distinct, we can ignore this distinction.

- 1. $\pi(u_1), \ldots, \pi(u_{d-1}) \ge r;$
- 2. wCut_r $(T_{w_1}), \ldots,$ wCut_r (T_{w_d}) happen;
- 3. $\operatorname{Cut}_r(T_v)$ does not happen.

These events are independent. Their probabilities are

- 1. $\Pr[\pi(u_i) \ge r] = 1 r$ for all $1 \le i \le d 1$; this holds since v is maximal in W, and thus $u_i \notin W$ and $\delta_{u_i} = 0$;
- 2. $\Pr[\operatorname{wCut}_r(T_{w_i})] \leq P_r = \frac{r}{1-r}$ since $\Pr[\pi(l) < r] \leq r \delta_l \leq r$ for all labels in T, all nodes have distinct labels, and T_{w_i} is an infinite complete binary tree;
- 3. $\Pr[\neg \operatorname{Cut}_r(T_v)] = 1 \Pr[\operatorname{Cut}_r(T_v)] \le 1 Q_{r-\delta_{\max}}$. This holds since T_v is a complete infinite binary tree and $\Pr[\pi(w) < r] < r \delta_{\max}$ for all nodes w in T_v

Multiplying everything together, we see that

$$\begin{aligned} \Pr_{D_v}[\operatorname{Cut}_r(T)] &- \Pr_D[\operatorname{Cut}_r(T)] = \delta_v \Pr_{\tau \sim D}[v \text{ is pivotal under } \tau] \\ &= \delta_v (1-r)^{d-1} \cdot \left(\frac{r}{1-r}\right)^d \cdot (1-Q_{r-\delta_{\max}}) \\ &= \delta_v \frac{1-Q_{r-\delta_{\max}}}{1-r} r^d . \end{aligned}$$

This concludes the proof of Lemma 49.

Proposition C.3 (TwoCC-cleanup; Proposition 51, restated). If T' is as above and $\delta_u = -\epsilon \gamma_{\text{TwoCC}}(r)$, then $\Pr[\text{Cut}_r(T')] \ge \Pr[\text{Cut}_r(\text{CLEANSUBTREE}(u, T'))]$, for our specific choices of $\gamma_{\text{TwoCC}}(r) = 40r^{7/2}(1-2r)^2$, $\gamma(r) = r(1-2r)^{3/2}$, and $\epsilon \le 0.13$.

Proof. First of all, it is enough to prove that $\Pr[\operatorname{wCut}_r(T_u)] \geq \Pr[\operatorname{wCut}_r(T^\infty)] = P_r$. Since $\Pr[\operatorname{wCut}_r(T_u)] = (r - \delta_u) \vee \Pr[\operatorname{Cut}_r(T_u)]$ (for our "real-numbers" \vee meaning $a \vee b = a + b - ab$), it makes sense to bound $\Pr[\operatorname{Cut}_r(T_u)]$ from below. We can assume every node in T_u has exactly two children; if not so, keep on adding children. Also, we assume that $\delta_v \geq 0$ for all non-root nodes of T_u ; if not, we simply change δ_v to 0. These steps only decrease $\Pr[\operatorname{Cut}_r(T_u)]$. Writing d(v) to denote the depth of v in T_u , we apply Corollary 50 and get

$$\Pr[\operatorname{Cut}_{r}(T_{u})] \geq Q_{r} - 1.02 \frac{1 - Q_{r}}{1 - r} \cdot \sum_{v \in V(T_{u}) \setminus \{u\}} \delta_{v} r^{d(v)} , \qquad (43)$$

Note that we proved Corollary 50 only for finite sets W; the sum in (43) is infinite but only finitely many nodes have $\delta_v \neq 0$. Let us work on the sum in (43). We partition $V(T_u) \setminus \{u\}$ into $A \uplus B$ as follows: consider a node v; if $\delta_v > 0$ then v is an "old" node, already existing in T_x , and therefore it has a label $b = \text{varlabel}_{T_x}(v) \in V$. If $\{b, x\} \in H_{\text{low}}$ (equivalently, in our notation, if $v \triangleleft \text{root}(T_x)$, put it into A; otherwise (including the case $\delta_v = 0$), put it into B. In words, A contains all proper descendants v of u that have a bias $\delta_v > 0$ because they (more properly: their original labels) are neighbors with x in H_{low} . Since x has at most two neighbors x_1, x_2 in H_{low} , the set A can be partitioned into two antichains $A = A_1 \uplus A_2$. Note that $\delta_v \leq \delta_{\text{root}} + \delta_{\text{non-root}}$ for all $v \in A$, and $\delta_v \leq 2 \delta_{\text{non-root}}$ for all $v \in B$, and therefore

$$\sum_{v \in V(T_u) \setminus \{u\}} \delta_v r^{d(v)} \leq (\delta_{\text{root}} + \delta_{\text{non-root}}) \sum_{a \in A} r^{d(v)} + 2 \,\delta_{\text{non-root}} \sum_{b \in B} r^{d(v)}$$

$$= (\delta_{\text{root}} + \delta_{\text{non-root}}) \sum_{a \in A} r^{d(v)} + 2 \,\delta_{\text{non-root}} \sum_{b \in A \cup B} r^{d(v)} - 2 \,\delta_{\text{non-root}} \sum_{a \in A} r^{d(v)}$$

$$= (\delta_{\text{root}} - \delta_{\text{non-root}}) \sum_{a \in A} r^{d(v)} + 2 \,\delta_{\text{non-root}} \sum_{b \in A \cup B} r^{d(v)}$$

$$\leq \delta_{\text{root}} \sum_{a \in A} r^{d(v)} + 2 \,\delta_{\text{non-root}} \sum_{b \in A \cup B} r^{d(v)} \,.$$

We bound the second sum by

$$\sum_{b \in A \cup B} r^{d(v)} \le \sum_{d=1}^{\infty} 2^d r^d = \frac{2 \, r}{1 - 2 \, r} \ ,$$

since $A \cup B$ contains at most 2^h nodes at depth h, and none at depth 0 (remember that the root u of T_u is not in $A \cup B$). To bound the first one, we need the following simple proposition:

Proposition 77. Let T be a binary tree and A an antichain of nodes in T, and $r \leq 1/2$. $\sum_{v \in A} r^{d_T(v)} \leq 1$.

We apply this proposition to the antichains A_1 and A_2 and see that $\sum_{a \in A} r^{d(v)} \leq 2$. Putting everything together,

$$\begin{aligned} (43) &\geq Q_r - 1.02 \, \frac{1 - Q_r}{1 - r} \cdot \sum_{v \in V(T_u) \setminus \{u\}} \delta_v r^{d(v)} \\ &\geq Q_r - 1.02 \, \frac{1 - Q_r}{1 - r} \cdot 2 \, \delta_{\text{root}} \\ &\quad -1.02 \, \frac{1 - Q_r}{1 - r} \cdot 2 \, \delta_{\text{non-root}} \, \frac{2 \, r}{1 - 2 \, r} \\ &\geq Q_r - \frac{2.04}{(1 - r)^3} \, ((1 - 2 \, r) \delta_{\text{root}} + 2 \, r \delta_{\text{non-root}}) \qquad (\text{since } 1 - Q_r = \frac{1 - 2 \, r}{(1 - r)^2}) \\ &=: Q_r - \text{loss} \; . \end{aligned}$$

For the reader who is interested in technical details, the reason we used the fact that A is the union of two antichains is to preserve the 1 - 2r factor in front of δ_{root} ; this is important since δ_{root} behaves "less nicely" as $r \to 1/2$ than $\delta_{\text{non-root}}$. We can now bound $\Pr[\text{wCut}_r(T_u)]$:

$$\Pr[\operatorname{wCut}_r(T_u)] \ge (r - \delta_u) \lor (Q_r - \operatorname{loss})$$

= $r - \delta_u + (1 - r + \delta_u) (Q_r - \operatorname{loss})$
= $r - \delta_u + (1 - r)Q_r - (1 - r)\operatorname{loss} + \delta_u Q_r - \delta_u \operatorname{loss}$
 $\ge r - \delta_u + (1 - r)Q_r - (1 - r)\operatorname{loss} + \delta_u Q_r$ (since $\delta_u \le 0$)
= $P_r - \delta_u (1 - Q_r) - (1 - r)\operatorname{loss}$.

We have to show that this is at least $\Pr[\operatorname{wCut}_r(T^{\infty})] = P_r$, which holds if and only if

$$-\delta_u(1-Q_r) \ge (1-r) \text{loss} \iff$$

$$\epsilon \gamma_{\text{TwoCC}}(r) \frac{1-2r}{(1-r)^2} \ge (1-r) \frac{2.04}{(1-r)^3} \left((1-2r)\delta_{\text{root}} + 2r\delta_{\text{non-root}} \right) \iff$$

$$\epsilon \gamma_{\text{TwoCC}}(1-2r) \ge 2.04 \left((1-2r)\delta_{\text{root}} + 2r\delta_{\text{non-root}} \right) \iff$$

$$\epsilon \gamma_{\text{TwoCC}}(1-2r) \ge 2.04 \cdot 1.2 \epsilon \gamma(r) \left((1-2r) \max(0, -\phi(r)) + \frac{2r\gamma(r)}{1-r} \right)$$

by the definition of δ_{root} and $\delta_{\text{non-root}}$ in (13) and (14). Solving for γ_{TwoCC} , this holds if and only if

$$\gamma_{\text{TwoCC}}(r) \ge 2.04 \cdot 1.2 \, \gamma(r) \left(\max(0, -\phi(r)) + \frac{2 \, r \gamma(r)}{(1-r)(1-2 \, r)} \right) \,.$$
 (44)

We verify numerically that this holds for our specific choice of $\gamma(r) = r(1-2r)^{3/2}$ and $\gamma_{\text{TwoCC}}(r) = 40r^{7/2}(1-2r)^2$. This concludes the proof of Proposition 51.

Proposition C.4 (One-child cleanup; Proposition 52, restated). Suppose T' is as above and u has at most child. Furthermore, suppose that $\delta_u \leq \delta_{\max}$ for all nodes u, and $r(1-2r) \geq 2 \delta_{\max}$. Then $\Pr[\operatorname{Cut}_r(T')] \geq \Pr[\operatorname{Cut}_r(\operatorname{CLEANSUBTREE}(u, T'))]$.

Proof. As with the previous proposition, it suffices to prove $\Pr[\operatorname{wCut}_r(T_u)] \geq \Pr[\operatorname{wCut}_r(T^{\infty})] = P_r$. Suppose u has exactly one child (the case that it has none is even easier), and call it v. Observe that $\operatorname{wCut}_r(T_u) = [\pi(u) < r] \lor \operatorname{wCut}_r(T_v)$. Since $\Pr[\pi(l) < r] \ge r - \delta_{\max}$ for all labels in T_u , the following simple lower bound on $\Pr[\operatorname{wCut}_r(T_v)]$ holds:

$$\begin{aligned} \Pr[\operatorname{wCut}_r(T_v)] &\geq P_{r-\delta_{\max}} \\ &\geq P_r - \delta_{\max} P'_r \\ &= P_r - \frac{\delta_{\max}}{(1-r)^2} . \end{aligned}$$
(since P_r is convex)

Now we can bound $\Pr[\operatorname{wCut}_r(T_u)]$:

$$\begin{aligned} \Pr[\operatorname{wCut}_r(T_u)] &= \Pr[\pi(u) < r \lor \operatorname{wCut}_r(T_v)] \\ &\geq (r - \delta_{\max}) \lor \left(P_r - \frac{\delta_{\max}}{(1 - r)^2} \right) \\ &= r - \delta_{\max} + (1 - r - \delta_{\max}) \left(P_r - \frac{\delta_{\max}}{(1 - r)^2} \right) \\ &= r - \delta_{\max} + (1 - r) P_r - \delta_{\max} P_r - \frac{(1 - r)\delta_{\max}}{(1 - r)^2} + \frac{\delta_{\max}^2}{(1 - r)^2} \\ &\geq r - \delta_{\max} + (1 - r) P_r - \delta_{\max} P_r - \frac{(1 - r)\delta_{\max}}{(1 - r)^2} \\ &= 2r - \frac{2\delta_{\max}}{1 - r} \qquad (\text{ since } P_r = \frac{r}{1 - r}) \end{aligned}$$

We have to show that this is at least $\Pr[\operatorname{wCut}_r(T_\infty)] = P_r = \frac{r}{1-r}$. Indeed,

$$2r - \frac{2\delta_{\max}}{1-r} - \frac{r}{1-r} = \frac{2r(1-r) - 2\delta_{\max} - r}{1-r} = \frac{r(1-2r) - 2\delta_{\max}}{1-r}$$

This is non-negative if and only if $r(1-2r) \ge 2\delta_{\max}$, which holds by assumption. This proves Proposition 52.
Proposition C.5 (Proposition 54, restated). $\frac{1.02\,\delta_{\text{non-root}}}{1-2\,r}r^{d-1} \leq 4.896\,\epsilon\left(\frac{1}{2}\right)^d \frac{(d+1)^{d+1}}{(d+3)^{d+3}}.$

Proof. By the definition of $\delta_{\text{non-root}}$ in (14), we get

$$\frac{1.02 \,\delta_{\text{non-root}}}{1-2 \,r} r^{d-1} = \frac{1.02 \cdot 1.2 \,\epsilon \gamma^2(r)}{(1-2 \,r)(1-r)} r^{d-1}$$
$$= \frac{1.224 \,\epsilon r^2(1-2 \,r)^3}{(1-2 \,r)(1-r)} r^{d-1}$$
$$\leq 2.448 \,\epsilon r^{d+1}(1-2 \,r)^2 \qquad (\text{since } 1-r \geq 1/2)$$

By basic calculus, the function $r^{d+1}(1-2r)^2$ is maximized for $r = \frac{1}{2} \cdot \frac{d+1}{d+3}$, at which point it becomes

$$\left(\frac{1}{2}\right)^{d+1} \left(\frac{d+1}{d+3}\right)^{d+1} \left(\frac{2}{d+3}\right)^2$$
$$= 2 \left(\frac{1}{2}\right)^d \frac{(d+1)^{d+1}}{(d+3)^{d+3}} .$$

This proves the proposition.

Proposition C.6. $\sum_{d=1}^{\infty} \frac{(d+1)^{d+1}}{(d+3)^{d+3}} \le 0.0544.$

Proof. Write $a_d := \frac{(d+1)^{d+1}}{(d+3)^{d+3}}$ and observe that $a_d = \frac{1}{(d+1)^2} \left(1 - \frac{2}{d+3}\right)^{d+3} < \frac{e^{-2}}{(d+1)^2} =: b_d$. Therefore,

$$\sum_{d=1}^{\infty} a_d = \sum_{d=1}^{\infty} b_d + \sum_{d=1}^{\infty} (a_d - b_d)$$
$$= e^{-2} \left(\frac{\pi^2}{6} - 1\right) + \sum_{d=1}^{\infty} (a_d - b_d)$$
$$< e^{-2} \left(\frac{\pi^2}{6} - 1\right) + \sum_{d=1}^{N} (a_d - b_d)$$

for any finite N, since $a_d - b_d < 0$. Plugging in N = 100 and evaluating this numerically proves the proposition.

Proposition C.7 (Proposition 57, restated). The marginal distribution of $(\pi(y), \pi(z))$, conditioned on $\pi(x) = r$, is still $D_{\epsilon}^{\gamma, \Box}$:

Proof. It is enough to show that $\Pr[\pi(y) \in A_y \land \pi(z) \in A_z \mid \pi(x) = r]$ has the right distribution for all intervals A_y, A_z , since intervals generate the whole σ -algebra of measurable sets. We apply Theorem 40 with $K = \{y, z\}$, $I = V \setminus K$, $A_x := \{r\}$, and $A_v := [0, 1]$ for all variables $v \in V \setminus \{x, y, z\}$; the latter choice reflects the fact that no condition is imposed on $\pi(v)$. By Theorem 40 we see that

$$\Pr[\pi(y) \in A_y, \pi(z) \in A_z \mid \pi(x) = r] = \mu(A_y)\mu(A_z) \cdot \left(1 + \epsilon \frac{\sum_{\{u,v\} \in E \setminus E_I} T_u T_v}{1 + \epsilon \sum_{\{u,v\} \in E_I} T_u T_v}\right) = \mu(A_y)\mu(A_z) \cdot (1 + \epsilon T_y T_z) .$$

Why does this hold? Every edge $\{u, v\} \in E_I$ contains some $v \in V \setminus \{x, y, z\}$ with $A_v = [0, 1]$ and therefore $T_v = 0$. The denominator becomes 1. As for the numerator, note that an edge $\{u, v\} \in E \setminus E_I$ is either $\{y, z\}$; or it contains again some $v \notin \{x, y, z\}$. Thus, all terms except $T_y T_z$ vanish. We continue:

$$\mu(A_y)\mu(A_z) \cdot (1 + \epsilon T_y T_z) = \mu(A_y)\mu(A_z) \cdot \left(1 + \epsilon \frac{\int_{A_y} \phi(s) \, ds}{\mu(A_y)} \cdot \frac{\int (A_z)\phi(t) \, dt}{\mu(A_z)}\right)$$
$$= \mu(A_y)\mu(A_z) + \epsilon \int_{A_y} \int_{A_z} \phi(s)\phi(t) \, dt \, ds$$
$$= \int_{A_y} \int_{A_z} (1 + \epsilon \phi(s)\phi(t)) \, dt \, ds$$
$$= \Pr_{(\pi(y),\pi(z)) \sim D_{\epsilon}^{\gamma,\Box}} [\pi(y) \in A_y \land \pi(z) \in A_z] .$$

This proves the proposition.

Lemma C.8 (Bonus from only child; Lemma 62, restated). Let $d = d_{T_x}(u)$ be the distance from the root to u. Then $\Pr[\operatorname{Cut}_r(T_{i-1})] \ge \Pr[\operatorname{Cut}_r(T_i)] + \operatorname{OCB}(d, r)$.

Proof. Let u be the i^{th} node of B_1 in our enumeration. We assume that u has exactly one child; the case that u has no child (i.e., is an unsafe leaf) is even better for us. Let root $= u_0, u_1, \ldots, u_d = u$ be the path from the root to u. For $1 \leq j \leq d$, let let w_j be the child of u_{j-1} that is not u_j . Note that u_0, \ldots, u_{d-1} are canonical in T_x and thus have exactly two children in T_x and thus in T'_x and T_{i-1} , as well. By T_{w_j} we denote the subtree of T_{i-1} rooted at w_j . Finally, let v be the only child of u in T_{i-1} .



Let $\tau \sim D$ be a placement of all labels except those of T_u and $T_{u'}$. Under τ , the events $\operatorname{Cut}_r(T_{i-1})$ and $\operatorname{Cut}_r(T_i)$ either (1) both become Ω (they both happen, regardless of what happens below u) or (2) both become \emptyset (they both don't happen, regardless of what happens below u), or (3) $\operatorname{Cut}_r(T_{i-1})$ becomes $\operatorname{wCut}_r(T_u)$ and $\operatorname{Cut}_r(T_i)$ becomes $\operatorname{wCut}_r(T_u)$. Call τ good if (3) happens. Note that $\Pr[\pi(l) < r] \geq r - \delta_{\max} =: \tilde{r}$ for every

label l in T_u . For a good τ , we get

$$\Pr_{D}[\operatorname{Cut}_{r}(T_{i-1}) \mid \tau] = \Pr_{D}[\operatorname{wCut}_{r}(T_{u})] = \Pr[\pi(u) < r \lor \operatorname{wCut}_{r}(T_{v})]$$
$$\geq \tilde{r} \lor P_{\tilde{r}} = \tilde{r} + (1 - \tilde{r})\frac{\tilde{r}}{1 - \tilde{r}} = 2 \tilde{r}$$
$$\Pr_{D}[\operatorname{Cut}_{r}(T_{i}) \mid \tau] = \Pr_{D}[\operatorname{wCut}_{r}(T_{u}) = P_{r} = \frac{r}{1 - r}.$$

Let Δ denote the difference between those quantities:

$$\Pr_{D}[\operatorname{Cut}_{r}(T_{i-1}) \mid \tau] - \Pr_{D}[\operatorname{Cut}_{r}(T_{i}) \mid \tau] \ge 2\,\tilde{r} - \frac{r}{1-r} = \frac{r(1-2\,r) - 2\,\delta_{\max}(1-r)}{1-r} =: \Delta \;.$$

We conclude that

$$\Pr[\operatorname{Cut}_r(T_{i-1})] - \Pr[\operatorname{Cut}_r(T_i)] \ge \Pr[\tau \text{ is good}] \cdot \Delta$$

Second, we compute the probability that τ is good. Note that τ is good if and only if

- 1. $\pi(u_1), \ldots, \pi(u_{d-1}) \ge r;$
- 2. $\operatorname{wCut}_r(T_{w_1}), \ldots, \operatorname{wCut}_r(T_{w_d})$ happen.

The events under Point 1 are independent, and each happens with probability $1-r+\delta_{u_i} \geq 1-r$. Those under Point 2 might not be independent, as the label z might occur in many of them. Still, they are monotone boolean functions in the events $[\pi(l) < r]$ and thus are positively correlated, which follows easily from the FKG inequality. Therefore, it holds that

$$\Pr_{D}[\operatorname{wCut}_{r}(T_{w_{1}}), \dots, \Pr_{r}(T_{w_{d}})] \ge \prod_{i=1}^{d} \Pr_{D}[\operatorname{wCut}_{r}(T_{w_{i}})] \qquad \text{(by FKG)}$$
$$\ge (P_{\tilde{r}})^{d} = (P_{r-\delta_{\max}})^{d}$$
$$\ge (P_{r} - \delta_{\max}P_{r}')^{d} \qquad \text{(since } P_{r} \text{ is convex)}$$
$$= \left(\frac{r}{1-r} - \frac{\delta_{\max}}{(1-r)^{2}}\right)^{d}$$

and

$$\Pr[\tau \text{ is good}] \ge (1-r)^{d-1} \cdot \left(\frac{r}{1-r} - \frac{\delta_{\max}}{(1-r)^2}\right)^d$$

Thus,

$$\begin{aligned} \Pr[\operatorname{Cut}_r(T_{i-1})] &- \Pr[\operatorname{Cut}_r(T_i)] = \Pr[\tau \text{ is good}] \cdot \Delta \\ &\geq (1-r)^{d-1} \cdot \left(\frac{r}{1-r} - \frac{\delta_{\max}}{(1-r)^2}\right)^d \cdot \frac{r(1-2r) - 2\delta_{\max}(1-r)}{1-r} \\ &= \left(r - \frac{\delta_{\max}}{1-r}\right)^d \cdot \frac{r(1-2r) - 2\delta_{\max}(1-r)}{(1-r)^2} \\ &= \operatorname{OCB}(d,r) \end{aligned}$$
 (as defined in (21))

This proves Lemma 62.

Lemma C.9 (Bonus from multiple labels; Lemma 63, restated). Let $d = d_{T_x}(u)$ be the distance from the root to u. Then $\Pr[\operatorname{Cut}_r(T'_{i-1})] \ge \Pr[\operatorname{Cut}_r(T'_i)] + \operatorname{MLB}(d)$.

Proof. As before, let u be the $(i - |B_1|)^{\text{th}}$ element of B_z and root $= u_0, \ldots, u_d = u$ be the path from the root to u. For $1 \le i \le d$, let w_i be the child of u_{i-1} that is not u_i . Consider a placement τ to all labels of T'_i and T'_{i-1} except z and l. Under τ , the events $\operatorname{Cut}_r(T'_{i-1})$ and $\operatorname{Cut}_r(T'_i)$ either

- 1. both become Ω (both happen regardless of z and l) or
- 2. both become \emptyset (both do not happen regardless of z and l) or
- 3. both become $[\pi(z) < r]$, or
- 4. $\operatorname{Cut}_r(T'_i)$ becomes $[\pi(z) < r \land \pi(l) < r]$ and $\operatorname{Cut}_r(T'_{i-1})$ becomes $[\pi(z) < r]$.

Note that "5. $\operatorname{Cut}_r(T'_i)$ becomes $[\pi(z) < r \lor \pi(l) < r]$ " is impossible since $B_z \cup \{R\}$ is an antichain. Let us call τ good if Point 4 happens. If τ is not good, then $\operatorname{Pr}[\operatorname{Cut}_r(T'_{i-1}) \mid \tau] = \operatorname{Pr}[\operatorname{Cut}_r(T'_i) \mid \tau]$. If τ is good, then

$$\Pr[\operatorname{Cut}_{r}(T'_{i}) \mid \tau] - \Pr[\operatorname{Cut}_{r}(T'_{i-1}) \mid \tau] = \Pr[\pi(z) < r] - \Pr[\pi(z) < r \land \pi(l) < r] = r(1-r)$$

We call τ excellent if the following things happen:

- 1. $\neg \operatorname{Cut}_r(T'_{w_1})$ and $\neg \operatorname{Cut}_r(T'_{u_d})$ happen under τ ;
- 2. $\tau(u_1), \ldots, \tau(u_{d-1}) \ge r;$
- 3. for $2 \le i \le d$, the event wCut_r (T'_{w_i}) either happens under τ or becomes $[\pi(z) < r]$ under τ

Note that if τ is excellent, then $\operatorname{Cut}_r(T'_i)$ and $\operatorname{Cut}_r(T'_{i-1})$ happen if and only if $\pi(z) < r$ and $\pi(\operatorname{varlabel}(u_d) < r$. Since u_d has label z in T'_{i-1} but label l in T'_i , this means that τ is good. What is the probability that τ is excellent? Note that τ is independent on all labels, and all labels of T'_{i-1} , T'_i are distinct except z. Thus, the events described in Points 1–3 are independent, and therefore

- 1. $\Pr[\neg \operatorname{Cut}_r(T'_{w_1}) \land \neg \operatorname{Cut}_r(T'_{u_d})] \ge (1 Q_r)^2$, since $\Pr[\tau(a) < r] \le r$ for all labels a.
- 2. $\Pr[\tau(u_1), ..., \tau(u_{d-1})] \ge (1-r)^{d-1}$ for the same reason.
- 3. Evaluating the probability that wCut_r (T'_{w_i}) becomes Ω or $[\pi(z) < r]$ is slightly more subtle. Since $\pi(z)$ is independent of τ , we get

$$\begin{aligned} \Pr_{\tau,\pi(z)}[\operatorname{wCut}_r(T'_{w_i})] &= & \Pr_{\tau}[\operatorname{wCut}_r(T'_{w_i}) \text{ happens under } \tau] \\ &+ & \Pr_{\tau}[\operatorname{wCut}_r(T'_{w_i}) \text{ becomes } [\pi(z) < r] \text{ under } \tau] \cdot \Pr[\pi(z) < r] \\ &\leq & \Pr_{\tau}[\operatorname{wCut}_r(T'_{w_i}) \text{ happens under } \tau] \\ &+ & \Pr_{\tau}[\operatorname{wCut}_r(T'_{w_i}) \text{ becomes } [\pi(z) < r] \text{ under } \tau] .\end{aligned}$$

Thus, the probability that wCut_r(T'_{w_i}) happens becomes or $[\pi(z) < r]$ under τ is at least $\Pr_D[\text{wCut}_r(T'_{w_i})]$, which is at least $P_{r-\delta_{\max}} \ge P_r - \frac{\delta_{\max}}{(1-r)^2}$. We conclude that Point 3 happens with probability at least $\left(P_r - \frac{\delta_{\max}}{(1-r)^2}\right)^{d-1}$.

Multiplying all probabilities, we get

$$\Pr[\operatorname{Cut}_{r}(T_{i-1}')] - \Pr[\operatorname{Cut}_{r}(T_{i}')] \ge \Pr[\tau \text{ is excellent}] \cdot r(1-r)$$

$$\ge (1-Q_{r})^{2}(1-r)^{d-1} \left(\frac{r}{1-r} - \frac{\delta_{\max}}{(1-r)^{2}}\right)^{d-1} r(1-r)$$

$$= (1-Q_{r})^{2} \left(r - \frac{\delta_{\max}}{1-r}\right)^{d-1} r(1-r)$$

$$= \operatorname{MLB}(d, r) \qquad (\text{as defined in (22)})$$

This concludes the proof of Lemma 63.

Lemma C.10 (Lemma 64, restated). $\sum_{v \in B_1} \text{OCB}(d(v)) + \sum_{v \in B_2} \text{MLB}(d(v)) \ge 0.9 \text{ Thr.}$

Proof of Lemma 64. Recall that through the bijection Φ , every $v \in B_1 \cup B_z$ corresponds to some a node $v \in A'_z$ in T_y , and $d_{T_x}(v) = d_{T_y}(u) + 1$. If $u \in A'_z \cap A_z$ then $\Phi(u) \in B_z$. Otherwise, if $u \in A'_z \setminus A_z$ then $\Phi(u) \in B_1$. Abbreviating $d(u) := d_{T_y}(u)$ (since a node should "know in which tree it lives"), we get

$$\sum_{v \in B_1} \text{OCB}(d(v)) + \sum_{v \in B_z} \text{MLB}(d(v)) = \sum_{u \in A'_z \cap A_z} \text{MLB}(d(u) + 1) + \sum_{u \in A'_z \setminus A_z} \text{OCB}(d(u) + 1)$$

If $u \in A'_z \setminus A_z$ then u is an ancestor of several $w \in A_z$; let A^u_z be this set. Note that $(A'_z \cap A_z) \cup \bigcup_{u \in A'_z \setminus A_z} A^u_z$ is a partition of A_z . Also note that $OCB(d+1) \leq \frac{1}{2} OCB(d)$, which is obvious from the fact that $r \leq 1/2$ and the definition of OCB. The following proposition is easily proved by induction:

Proposition C.11. If $f : \mathbb{N}_0 \to \mathbb{R}_0^+$ is a function such that $f(d+1) \leq f(d)/2$ and A_u is an antichain of descendants of u in some binary tree, then $f(d(u)) \geq \sum_{w \in A_u} f(d(w))$.

With this proposition, we see that the above sum equals

$$\sum_{u \in A'_z \cap A_z} \operatorname{MLB}(d(u)+1) + \sum_{u \in A'_z \setminus A_z} \sum_{w \in A^u_z} \operatorname{OCB}(d(w)+1)$$

$$\geq \sum_{u \in A_z} \min \left(\operatorname{MLB}(d(u)+1), \operatorname{OCB}(d(u)+1) \right)$$
(45)

Proposition C.12. Provided that $\epsilon \leq 0.1$, it holds that

$$OCB(d) = \int_0^{1/2} OCB(d, r) \, dr \ge 0.88 \cdot \int_0^{1/2} \frac{r(1-2r)}{(1-r)^2} \cdot r^d \, dr =: OCB^*(d)$$
$$MLB(d) = \int_0^{1/2} MLB(d, r) \, dr \ge 0.9 \cdot \int_0^{1/2} \frac{(1-2r)^2}{(1-r)^3} \cdot r^d \, dr =: MLB^*(d)$$

Proof. As a reminder, we copy the definitions of OCB and MLB:

$$OCB(d,r) := \frac{r(1-2r) - 2\delta_{\max}(1-r)}{(1-r)^2} \cdot \left(r - \frac{\delta_{\max}}{1-r}\right)^d , \qquad \text{(as defined in (21))}$$
$$MLB(d,r) := \frac{(1-2r)^2r}{(1-r)^3} \cdot \left(r - \frac{\delta_{\max}}{1-r}\right)^{d-1} . \qquad \text{(as defined in (22))}$$

We start with the first integral. Referring to the definition of δ_{\max} in (16) and to our promise that $\epsilon \leq 0.1$, we can verify numerically that $\frac{r(1-2r)-2\delta_{\max}(1-r)}{(1-r)^2} \geq 0.95 \frac{r(1-2r)}{(1-r)^2}$. We set $\eta(r) := \frac{\delta_{\max}}{1-r}$ and $s(r) = r - \eta(r)$ and abbreviate $f(r) := \frac{r(1-2r)}{(1-r)^2}$. With this notation, we conclude that

$$OCB(d) = \int_0^{1/2} OCB(d,r) \, dr \ge 0.95 \, \int_0^{1/2} f(r)(s(r))^d \, dr \; .$$

Claim. $f(r) \ge 0.98 \cdot f(s(r)).$

This can easily be verified numerically. Using the claim, we see that

OCB(d)
$$\ge 0.95 \cdot 0.98 \int_0^{1/2} f(s(r))(s(r))^d dr$$

Claim. $s'(r) \leq 1.05$ for $r \in [0, 1/2]$, provided $\epsilon \leq 0.13$.

This can also be checked numerically. Technically, since δ_{\max} is defined by a maximum, the function s(r) fails to be differentiable at $r = \frac{5-\sqrt{13}}{6} \approx 0.2324$. However, this is a single point and thus does not affect the integral. The above integral is at least

$$\frac{0.95 \cdot 0.98}{1.05} \int_0^{1/2} f(s(r))(s(r))^d s'(r) \, dr = \frac{0.95 \cdot 0.98}{1.05} \int_0^{1/2} f(s) s^d \, ds \ge \text{OCB}^*(d) \; .$$

This proves the bound for the first integral.

For the second integral, we use the same $\eta(r), s(r)$. Writing $g(r) = \frac{(1-2r)^2}{(1-r)^3}$, we get

$$\begin{split} \int_{0}^{1/2} \frac{(1-2r)^{2}r}{(1-r)^{3}} \cdot \left(r - \frac{\delta_{\max}}{1-r}\right)^{d-1} dr &\geq \int_{0}^{1/2} g(r) \cdot (s(r))^{d} dr \\ &\geq \int_{0}^{1/2} g(r) \cdot (s(r))^{d} s'(r) \frac{1}{s'(r)} dr \\ &\geq \frac{1}{1.05} \cdot \int_{0}^{1/2} g(r) \cdot (s(r))^{d} s'(r) dr \;, \end{split}$$

using the above claim that $s'(r) \leq 1.05$. We need to bound g(r) from below.

Claim. $g(r) \ge 0.945 \cdot g(s(r)).$

Altogether, the integral is at least

$$\frac{0.945}{1.05} \cdot \int_0^{1/2} g(s(r)) \cdot (s(r))^d s'(r) \, dr = \frac{0.945}{1.05} \cdot \int_0^{1/2} g(s) s^d \, ds$$
$$\ge \mathrm{MLB}^*(d) \; .$$

This shows the bound for the second integral, and proves the proposition.

Proposition C.13. The following inequalities hold:

1. For $d \ge 5$, it holds that $OCB^*(d) \ge MLB^*(d)$.

2. For $d \le 4$, $OCB^*(d)$, $MLB^*(d) \ge \frac{1}{1150} \ge THR$.

Proof. The second point can simply be verified numerically. The first one can be verified up for any fixed value of d. Beyond that, we need to proceed symbolically. As in the previous proof, let $f(r) := \frac{r(1-2r)}{(1-r)^2}$ and $g(r) := \frac{(1-2r)^2}{(1-r)^3}$. Note that $0 \le f(r), g(r) \le 1$. Furthermore, it is straightforward to check that $g(r) \le \frac{1}{2} f(r)$ for all $\theta := 0.45 \le r \le 1/2$. Intuitively, for large d, the factor r^d in the integrand pulls the bulk of the mass of the integral $\int_0^{1/2} g(r)r^d dr$ beyond θ , where $g(r) \le f(r)/2$ holds. More formally, observe that

$$\int_0^\theta g(r)r^d \le \int_0^\theta r^d = \frac{\theta^{d+1}}{d+1}$$

Set $r_{\min} := \frac{1}{2} \frac{d-1}{d+1}$ and $r_{\max} := \frac{1}{2} \frac{d+1}{d+3}$. For $r_{\min} \le r \le r_{\max}$, it holds that

$$g(r)r^d \ge (1-2r)^2 r^d \ge \left(\frac{2}{d+3}\right)^2 \left(\frac{d-1}{d+1}\right)^d 2^{-d}$$

The factor $\left(\frac{d-1}{d+1}\right)^d$ increases in d and is at least 1/10 for $d \ge 2$. Thus, the above term is at least

$$\frac{4}{10} \frac{1}{(d+3)^2} 2^{-d}$$

For all $d \geq 19$, it holds that $r_{\min} \geq \theta$, and therefore

$$\begin{split} \int_{\theta}^{1/2} g(r) r^d \, dr &\geq \int_{r_{\min}}^{r_{\max}} g(r) r^d \, dr \\ &\geq \left(r_{\max} - r_{\min} \right) \frac{2}{5} \frac{1}{(d+3)^2} 2^{-d} \\ &= \frac{2}{(d+3)(d+1)} \frac{2}{5} \frac{1}{(d+3)^2} 2^{-d} \\ &\geq \frac{4}{5} \frac{2^{-d}}{(d+1)(d+3)^3} \; . \end{split}$$

Comparing $\int_0^\theta g(r) r^d dr$ to $\int_\theta^{1/2} g(r) r^d dr$, we see that

$$\frac{\int_0^\theta g(r)r^d \, dr}{\int_{\theta}^{1/2} g(r)r^d \, dr} \le \frac{\frac{\theta^{d+1}}{d+1}}{\frac{4}{5} \frac{2^{-d}}{(d+1)(d+3)^3}} \le (2\,\theta)^d \cdot \frac{5\,\theta(d+3)^3}{4} \ .$$

Since $2\theta < 1/2$ and does not depend on d, the above expression converges to 0 and is at most 1/10 for all $d \ge 162$. Thus,

$$\begin{split} \text{MLB}^*(d) &= 0.9 \, \int_0^{1/2} g(r) r^d \, dr = 0.9 \left(\int_0^\theta g(r) r^d \, dr + \int_\theta^{1/2} g(r) r^d \, dr \right) \\ &\leq 0.9 \cdot \frac{11}{10} \cdot \int_\theta^{1/2} g(r) r^d \, dr \\ &\leq \frac{1}{2 \cdot 0.88} \int_\theta^{1/2} 0.88 \, f(r) r^d \, dr \qquad \text{(by choice of } \theta) \\ &\leq \frac{1}{2 \cdot 0.88} \int_0^{1/2} 0.88 \, f(r) r^d \, dr \qquad \text{(integrand is non-negative)} \\ &= \frac{1}{2 \cdot 0.88} \, \text{OCB}^*(d) \, . \end{split}$$

This shows that $OCB^*(d) \ge MLB^*(d)$ for all $d \ge 162$. For all smaller values of d, we verify the inequality by numerical computation. This concludes the proof.

There are two cases. If A_z contains some u with $d_{T_y}(u) \leq 3$ then $(45) \geq \min(\text{MLB}(4), \text{OCB}(4)) \geq$ THR, and the statement of Lemma 64 holds. Otherwise,

$$\begin{split} (45) &\geq \sum_{u \in A_z} \min \left(\text{MLB}(d(u) + 1), \text{OCB}(d(u) + 1) \right) \\ &\geq \sum_{u \in A_z} \min \left(\text{MLB}^*(d(u) + 1), \text{OCB}^*(d(u) + 1) \right) \\ &\geq \sum_{u \in A_z} \text{MLB}^*(d(u) + 1) \qquad (\text{since } d(u) + 1 \ge 5) \\ &= 0.9 \cdot \sum_{u \in A_z} \int_0^{1/2} \frac{(1 - 2r)^2}{(1 - r)^3} \cdot r^{d(u) + 1} dr \\ &= 0.9 \cdot \text{LABELDENSITY}(z, T_y) \;, \end{split}$$

for LABELDENSITY as defined in (10), since A_z is by definition the set of canonical nodes in T_y that have label z. By assumption, $\{y, z\} \in H_{\text{high}}$ and thus the label density is at least THR. This concludes the proof of Lemma 64.

D Proofs from Section 8

Lemma D.1 (TwoCC-cleanup in the irregular case; Lemma 72, restated). Let T be the (by now) infinite tree described above. Suppose u is a node in T with $\gamma_u = \gamma_{\text{TwoCC}}$. Then $\Pr[\text{Cut}_r(T)] \ge \Pr[\text{Cut}_r(\text{CLEANSUBTREE}(u,T)), \text{ provided that } \gamma_{\text{TwoCC}}(r) \ge \frac{2r\gamma_{\text{ID}_{0,1}}}{1-2r}.$

Proof. The proof is similar to (but simpler than) that of Lemma 51. Let T_u be the subtree of T rooted at u. It suffices to show that $\Pr[\operatorname{wCut}_r(T_u)] \ge \Pr[\operatorname{wCut}_r(T^{\infty})] = P_r$.

First, consider any proper descendant v of u. Since all labels in T are distinct, we take the liberty of writing $\pi(v)$ instead of $\pi(\text{varlabel}(v))$. A brief glance at the definition of D^B reveals that $\Pr[\pi(v) < r] \ge r - \epsilon \gamma_{\text{ID}_{0,1}}(r) =: r - \delta$, and therefore

$$\begin{aligned} \Pr[\operatorname{Cut}_r(T_u)] &\geq Q_{r-\delta} & \text{(pessimistically)} \\ &\geq Q_r - \delta Q'_r & \text{(since } Q_r \text{ is convex)} \\ &= Q_r - \frac{2 r \delta}{(1-r)^3} . \end{aligned}$$

By assumption, we have $\Pr[\pi(u) < r] = r + \epsilon \gamma_{\text{TwoCC}}(r) =: r + \delta_u$, and therefore

$$\begin{aligned} \Pr[\operatorname{wCut}_r(T_u)] &= \Pr[\pi(u) < r] \lor \Pr[\operatorname{Cut}_r(T_u)] \\ &\geq (r + \delta_u) \lor \left(Q_r - \frac{2 r \delta}{(1 - r)^3}\right) \\ &= r + \delta_u + (1 - r - \delta_u) \left(Q_r - \frac{2 r \delta}{(1 - r)^3}\right) \\ &= r + \delta_u + (1 - r)Q_r - \delta_u Q_r - \frac{2 r \delta}{(1 - r)^2} + \frac{2 r \delta \delta_u}{(1 - r)^3} \\ &\geq r + \delta_u + (1 - r)Q_r - \delta_u Q_r - \frac{2 r \delta}{(1 - r)^2} \\ &= P_r + \delta_u (1 - Q_r) - \frac{2 r \delta}{(1 - r)^2} \end{aligned}$$

We have to show that this is P_r , which happens if and only if

$$\begin{split} \delta_u(1-Q_r) &\geq \frac{2\,r\delta}{(1-r)^2} &\iff \\ \frac{\delta_u(1-2\,r)}{(1-r)^2} &\geq \frac{2\,r\delta}{(1-r)^2} &\iff \\ \delta_u &\geq \frac{2\,r\delta}{1-2\,r} &\iff \\ \gamma_{\mathrm{TwoCC}}(r) &\geq \frac{2\,r\gamma_{\mathrm{ID}_{0,1}}}{1-2\,r} \ . \end{split}$$

This concludes the proof of Lemma 72.

ISSN 1433-8092

https://eccc.weizmann.ac.il