

Linear Variance of the Random 3-SAT Hitting Time

OpenAI

October 5, 2026

Abstract

For random 3-SAT on n Boolean variables, with independent uniformly signed clauses on three distinct variables sampled with replacement, we prove that the first unsatisfiable prefix has variance $\Theta(n)$. The upper bound removes the logarithmic loss in the earlier variance estimate; the matching lower bound follows from Wilson's transition-width theorem.

1 Introduction

Fix an integer $n \geq 3$. A *proper 3-clause* is a disjunction of three literals on distinct variables among $x_1, \dots, x_n$. Let $C_1, C_2, \dots$ be independent clauses, each uniform among the $8\binom{n}{3}$ proper 3-clauses. Thus the signs are fair and independent, and whole clauses may repeat. Write

$$F_m = \bigwedge_{j=1}^m C_j, \quad H_n = \min\{m \geq 1 : F_m \text{ is unsatisfiable}\},$$

where F_0 is the empty, satisfiable formula. The random variable H_n records the location of the satisfiability transition in the clause process.

Remark (notation). $n \geq 3$ is the integer number of variables in the proper 3-clause process. In Lemma 3.1, Lemma 3.2, Lemma 4.1, Corollary 4.2, Proposition 4.3, Lemma 4.4: The variable count satisfies $n \geq 6$. j is an integer index: clause indices in C_j start at 1, whereas the number of auxiliary proper 2-clauses in $q_j(S)$ starts at 0; other summation ranges are those displayed locally. m is a nonnegative integer indexing prefixes of the clause process. In Lemma 3.2, Lemma 4.1: The time index satisfies $i \leq m < L$. In Proposition 4.3: In the proof's sums over prefix times, m ranges from i to $L - 1$. C has a local meaning: an absolute finite constant in the variance bounds or a clause in the survivor-set construction. In Theorem 1.1, Corollary 5.2: C is an absolute finite constant satisfying $\text{Var}(H_n) \leq Cn$ for every integer $n \geq 3$. In Lemma 2.2: C in the potential increment denotes an independent uniform proper 3-clause on the u variables.

Theorem 1.1. *There is an absolute constant $C < \infty$ such that*

$$\text{Var}(H_n) \leq Cn \quad (n \geq 3).$$

There are also absolute constants $c > 0$ and n_0 such that $\text{Var}(H_n) \geq cn$ for $n \geq n_0$. In particular, $\text{Var}(H_n) = \Theta(n)$ as $n \rightarrow \infty$.

Remark (notation). n_0 is an absolute deterministic threshold, independent of n , such that $\text{Var}(H_n) \geq cn$ for every admissible integer $n \geq n_0$.

The new assertion is the upper bound. We give its proof in full. The lower bound follows from Wilson's quantitative separation of satisfiability levels [7, Corollary 4]; Section 5 includes the short deduction, with the integer endpoints treated explicitly.

Context and contribution

Friedgut proved sharpness of the random k -SAT threshold for every fixed $k \geq 2$ [5, Theorem 1.3]. Quantifying the width of that transition is a finer problem. Wilson showed that, for $k \geq 3$, a fixed decrease in the satisfiability probability requires at least a constant times $\sqrt{n}$ additional clauses [7]. More recently, Carenini obtained an $O(n/\log n)$ central clause-window bound for fixed $k \geq 3$, including independently sampled proper constraints [3, Corollaries 7.10–7.11]. Her subsequent paper [2, Theorems 1.1 and 1.3, Corollary 1.2], made public on October 5, 2026, proves for every fixed $k \geq 3$ a central window bound $O_{k,\eta}(n^{1/2+1/k})$ at each fixed $0 < \eta < 1/2$ and a hitting-time variance bound $O_k(n^{1+2/k})$. Together with the Abbe–Montanari criterion, her result establishes the limiting satisfiability threshold for every fixed $k \geq 3$. We credit Carenini with priority for the resolution of the satisfiability conjecture. The present paper concerns the sharper fluctuation question for three-literal clauses.

Remark (notation). $\eta \in (0, 1/2)$ is a fixed deterministic probability level, independent of n ; constants in $\Theta_\eta(\sqrt{n})$ may depend on η .

A bound on a fixed central window alone does not control the variance, which also depends on the tails. In the other direction, Theorem 1.1 and Wilson’s lower bound show that the clause window over which satisfiability drops from $1 - \eta$ to η has order $\sqrt{n}$ for every fixed $0 < \eta < 1/2$; see Corollary 5.2.

The companion preprint *Variance of the Random k -SAT Hitting Time* [6] develops a clause-omission and root-test framework. It proves variance $\Theta_k(n)$ for every fixed $k \geq 4$ and an upper bound $O(n \log n)$ for $k = 3$. Our argument uses that framework, which we reprove here, and replaces its critical occupation estimate in the three-literal case. The companion has broader general- k and capped formulations; the present paper concerns the uncapped proper 3-clause process defined above.

The new ingredient is a family of potentials for an arbitrary set S of Boolean assignments. Let $q_j(S)$ be the probability that j independent random proper 2-clauses leave no assignment in S , and put

$$P_d(S) = \sum_{j=0}^{d-1} q_j(S), \quad d \geq 1.$$

The potential increases when S shrinks and is bounded by d . Section 2 proves that $q_d(S)^{3/2}$, for nonempty S , is controlled by the expected increment of P_d under a random 3-clause, up to a small error. Summing those increments gives a direct bound on the accumulated values of $q_d^{3/2}$. This is the step that removes the logarithmic loss.

Remark (notation). S denotes a subset of $\{0, 1\}^u$ to which the clause-survival and killing functionals are applied. In Lemma 2.1: S is nonempty. In Lemma 3.2, Lemma 4.1, Corollary 4.2, Proposition 4.3: The assignment sets used here are S_m , for $i \leq m \leq L$: the nonroot assignments that, together with the root assignment w falsifying C_i , satisfy $B_m^{(i)}$; these sets are $\mathcal{G}_m$ -measurable and decrease with m . For every integer $j \geq 0$ and $S \subseteq \{0, 1\}^u$, $q_j(S)$ is the probability that j independent random proper 2-clauses kill S , with $q_j(\emptyset) = 1$ and $q_0(S) = 0$ for $S \neq \emptyset$. d is a positive integer used as the clause count in $q_d(S)$ and the length of the sum $P_d(S) = \sum_{j=0}^{d-1} q_j(S)$. In Lemma 2.2: $d \geq 1$ is an arbitrary fixed integer. In Lemma 3.2, Lemma 4.1, Corollary 4.2, Proposition 4.3, Lemma 4.4: $d = \max\{1, Z\}$, where Z counts indices $j \leq L$, excluding i , whose clauses meet at least one root; in particular, d is $\mathcal{E}$ -measurable.

Here is how the potential enters the variance proof. For $n \geq 6$, cap H_n at $10n$ and let D_i be the delay caused by omitting the clause at index i , without renumbering later clauses. The coordinate-

omission form of the Efron–Stein inequality bounds the capped variance by $\sum_i \mathbb{E} D_i^2$ [4, 1]. Fixing the omitted clause singles out its three variables as roots. If that clause is pivotal, every remaining solution has the unique root assignment that falsifies it. Flipping each root gives an independent family of necessary tests on the nonroot solution set, except when another clause meets two or more roots. Section 3 constructs this exposure and proves the test bound.

Remark (notation). i is an integer index with $1 \leq i \leq L$ specifying the omitted clause C_i ; omission retains the original indices of all other clauses. For $n \geq 6$, $L = 10n$ and $1 \leq i \leq L$, $D_i = T^{(i)} - T \geq 0$ is the delay obtained by omitting C_i while retaining all original clause indices; it is a function of the first L clauses. In Lemma 3.2, Lemma 4.1, Corollary 4.2, Proposition 4.3: $D_i = \sum_{m=i}^{L-1} \mathbf{1}_{A_m}$, where $A_m = \{T \leq m < T^{(i)}\}$.

Let d be the larger of 1 and the number of clauses other than C_i among the first $10n$ that meet a root. At a given index, let S be the nonroot solution set of the omitted-clause prefix with the roots fixed to their falsifying assignment, and write $p = q_d(S)$. With no collisions, the three test families each have at most d clauses and give a pivotality bound p^3 . For $p > 0$, the potential gives a conditional remaining-duration bound $O(d^{3/2}p^{-3/2})$ and bounds the conditional expected sum of $p^{3/2}$, while S is nonempty, by $O(d^{3/2})$. The identity $3 - 3/2 = 3/2$ thus yields a conditional squared-delay bound $O(d^3)$. The root-incidence count d has bounded moments uniformly in n . Section 4 carries out this calculation and averages the rare collision cases, obtaining $\mathbb{E} D_i^2 = O(1)$ uniformly in i . Finally, Section 5 removes the cap using an elementary first-moment tail bound and proves the matching lower bound.

Remark (notation). In the introductory killing-probability shorthand $p = q_d(S)$, whereas in $r_n(p)$ the symbol p denotes a deterministic probability level satisfying $0 < p < 1$.

Throughout, a random proper a -clause on a specified variable set is uniform over the clauses using a distinct variables and independent fair signs. Independent collections are sampled with replacement. All unspecified constants are absolute.

Remark (notation). a is a positive deterministic parameter whose meaning is local to its occurrence. In Lemma 2.1, Lemma 2.2, Lemma 3.2, Lemma 4.1, Corollary 4.2, Proposition 4.3: In the short-clause notation, $a \in \{2, 3\}$ is the number of distinct variables in a proper clause. In Theorem 5.1: $a(p_1, p_2) > 0$ is a constant depending on the fixed levels $1 > p_1 > p_2 > 0$ and independent of n, m_1, m_2 .

2 A potential for killing assignment sets

We first work on $u \geq 3$ variables, independently of the clause process. The aim is to bound a power of a short-clause killing probability by the drift of a bounded, monotone potential. For $S \subseteq \{0, 1\}^u$ and a clause C , let

$$S[C] = \{z \in S : z \models C\}.$$

For $j \geq 0$, define $q_j(S)$ as the probability that j independent random proper 2-clauses *kill* S , meaning that no member of S satisfies them all. In particular,

$$q_j(\emptyset) = 1 \quad (j \geq 0), \quad q_0(S) = 0 \quad (S \neq \emptyset).$$

The quantity $q_j(S)$ is nondecreasing both in j and when S shrinks. Set $\beta = 3/2$ for the rest of the paper.

Remark (notation). $u \geq 3$ is the integer number of coordinates of the assignment space on which the clause-survival and killing functionals are evaluated. In Lemma 3.2, Lemma 4.1, Corollary 4.2, Proposition 4.3: $u = n - 3$ is the number of nonroot variables, with $n \geq 6$.

Lemma 2.1 (One-clause comparison). *For a nonempty assignment set S , let $h_a(S)$ be the probability that one random proper a -clause kills S , for $a = 2, 3$. Then*

$$h_2(S)^\beta \leq 3h_3(S) + u^{-3}. \quad (2.1)$$

Proof. A coordinate is *frozen* if it is constant throughout S . Let b be the number of frozen coordinates. A clause kills all of S exactly when all its variables are frozen and each literal has the falsifying sign. Consequently,

$$h_a(S) = \frac{(b)_a}{2^a(u)_a}, \quad a = 2, 3,$$

where $(t)_a = t(t-1)\cdots(t-a+1)$, with value zero for integers $0 \leq t < a$. Since $b \leq u$,

$$\sqrt{h_2(S)} \leq \frac{b}{2u}.$$

For $b \geq 3$ we also have

$$\frac{h_3(S)}{h_2(S)} = \frac{b-2}{2(u-2)} \geq \frac{b}{6u},$$

which proves $h_2(S)^{3/2} \leq 3h_3(S)$ in this case. For $b \leq 1$ the left side is zero. For $b = 2$ it is $[2u(u-1)]^{-3/2} \leq u^{-3}$, proving (2.1). $\square$

Remark (notation). t is a deterministic real parameter with a meaning local to its occurrence. In Lemma 2.1: t is the argument of $(t)_a = t(t-1)\cdots(t-a+1)$; the evaluations used are $t = b$ and $t = u$. In Corollary 5.2: $t = (2Cn/\eta)^{1/2}$, with C the upper-bound constant from Theorem 1.1; this deviation radius varies with n .

The error term accounts for sets with exactly two frozen coordinates: a 2-clause can kill such a set, whereas a proper 3-clause cannot. The following lemma turns the one-clause comparison into the potential estimate needed along a decreasing sequence of assignment sets.

Lemma 2.2 (Potential drift). *For every integer $d \geq 1$, put*

$$P_d(S) = \sum_{j=0}^{d-1} q_j(S).$$

Then $0 \leq P_d(S) \leq d$, and P_d is nondecreasing when S shrinks. If C is an independent random proper 3-clause on the u variables, then

$$\mathbf{1}_{\{S \neq \emptyset\}} q_d(S)^\beta \leq d^{\beta-1} \left(3\mathbb{E}_C[P_d(S[C]) - P_d(S)] + du^{-3} \right). \quad (2.2)$$

Proof. Fix $j \geq 0$, and let Y be the subset of S surviving j independent random proper 2-clauses. Adding one more such clause gives

$$q_{j+1}(S) - q_j(S) = \mathbb{E}[\mathbf{1}_{\{Y \neq \emptyset\}} h_2(Y)]. \quad (2.3)$$

Adding instead the independent 3-clause C gives

$$\mathbb{E}_C[q_j(S[C]) - q_j(S)] = \mathbb{E}[\mathbf{1}_{\{Y \neq \emptyset\}} h_3(Y)]. \quad (2.4)$$

Indeed, in either order of adding the clauses, the increase in killing probability is precisely the probability that the old surviving set is nonempty and the final clause kills it. Products with

the indicators in these formulas are defined to be zero when Y is empty. Jensen's inequality and Lemma 2.1 imply

$$(q_{j+1}(S) - q_j(S))^\beta \leq 3\mathbb{E}_C[q_j(S[C]) - q_j(S)] + u^{-3}.$$

For nonempty S , the nonnegative successive differences telescope to $q_d(S)$, because $q_0(S) = 0$. Thus the power-sum inequality yields

$$q_d(S)^\beta \leq d^{\beta-1} \sum_{j=0}^{d-1} (q_{j+1}(S) - q_j(S))^\beta,$$

and summing the preceding bound proves (2.2). For empty S , its left side is zero and its right side is nonnegative. The bounds and monotonicity of P_d follow directly from those of q_j . $\square$

Remark (notation). $Y \subseteq S$ is the survivor set obtained by imposing the auxiliary proper 2-clauses. In Lemma 2.2: In the proof, j is a fixed nonnegative integer and Y consists of exactly those $z \in S$ satisfying all j auxiliary clauses; thus $Y = S$ when $j = 0$, and Y may be empty.

3 Clause deletion and independent root tests

We now reduce the variance bound to a uniform squared-delay estimate. The root exposure and root-flip tests in this section follow the clause-omission framework of [6, Section 3]; all details needed here are included.

Assume $n \geq 6$, set $L = 10n$, and let $T = \min(H_n, L)$. For $1 \leq i \leq L$, omit C_i while retaining the original indices:

$$B_m^{(i)} = \bigwedge_{\substack{1 \leq j \leq m \\ j \neq i}} C_j.$$

Let $T^{(i)}$ be the first index at which this omitted-clause process is unsatisfiable, capped at L , and set $D_i = T^{(i)} - T \geq 0$. Both capped variables are functions of the first L clauses, and $T^{(i)}$ does not depend on C_i .

Remark (notation). $T = \min(H_n, L)$ is defined for $n \geq 6$, with $L = 10n$, and is a function of the first L clauses. For $n \geq 6$ and $1 \leq i \leq L = 10n$, $T^{(i)}$ is the first unsatisfiable original index of $B_m^{(i)}$, capped at L ; it is a function of $(C_j)_{1 \leq j \leq L, j \neq i}$ and does not depend on C_i .

Lemma 3.1 (Coordinate omission). *The capped hitting time satisfies*

$$\text{Var}(T) \leq \sum_{i=1}^L \mathbb{E} D_i^2. \quad (3.1)$$

Proof. This is the coordinate-omission form of the Efron–Stein inequality; see [1, Section 3.2, Equation (3.6)]. For completeness, write $\mathcal{F}_i = \sigma(C_1, \dots, C_i)$ and $\Delta_i X = \mathbb{E}[X \mid \mathcal{F}_i] - \mathbb{E}[X \mid \mathcal{F}_{i-1}]$. Independence and the omission of C_i give $\Delta_i T^{(i)} = 0$, so $\Delta_i T = -\Delta_i D_i$. Orthogonality of conditional-expectation projections implies $\mathbb{E}(\Delta_i D_i)^2 \leq \mathbb{E} D_i^2$. Summing the orthogonal martingale differences of T proves (3.1). $\square$

Remark (notation). For $0 \leq i \leq L$, $\mathcal{F}_i = \sigma(C_1, \dots, C_i)$, with $\mathcal{F}_0$ the trivial sigma-field.

3.1 The exposure and its two filtrations

Fix i and condition on the value of C_i until further notice. All probabilities and expectations in this section and in Section 4, until the final averaging step, are in this conditional law. Write $B_m = B_m^{(i)}$. Call the three variables of C_i the *roots*, denote their set by R , and let $w \in \{0, 1\}^R$ be the unique root assignment falsifying C_i . There are $u = n - 3 \geq 3$ nonroot variables.

For $i \leq m < L$, define

$$A_m = \{B_m \text{ is satisfiable and } B_m \wedge C_i \text{ is unsatisfiable}\} = \{T \leq m < T^{(i)}\}. \quad (3.2)$$

The processes agree before i , so

$$D_i = \sum_{m=i}^{L-1} \mathbf{1}_{A_m}. \quad (3.3)$$

For $i \leq m \leq L$, let $S_m \subseteq \{0, 1\}^u$ consist of the nonroot assignments that, together with w , satisfy B_m . The sets S_m decrease with m . On A_m , the set S_m is nonempty, and every solution of B_m has root assignment w .

Remark (notation). A_m is defined for integer indices $i \leq m < L$ and is $\mathcal{H}_m$ -measurable. For $i \leq m \leq L$, $S_m \subseteq \{0, 1\}^u$, with $u = n - 3$, is the $\mathcal{G}_m$ -measurable set of nonroot assignments that together with w satisfy B_m , and these sets decrease with m .

We will retain two independent sources of randomness: clauses driving the evolution of S_m , and clauses testing possible root flips. For every index $j \leq L$ other than i , first reveal the subset of R met by C_j and all signs on that subset. Classify the index as follows:

- An *ordinary* clause meets no root; leave all its contents hidden.
- A *test* clause meets exactly one root, whose literal is true under w ; leave its two nonroot literals hidden.
- For every other clause, reveal all its remaining contents.

Let $\mathcal{E}$ be the sigma-field of this exposure, including the information at future indices. Let Z count the exposed indices meeting at least one root, and let V count those meeting at least two roots. The latter clauses are called *collisions*. Set $d = \max\{1, Z\}$. The counts Z, V, d and the complete schedule of clause types are $\mathcal{E}$ -measurable.

Remark (notation). Z is the $\mathcal{E}$ -measurable count $|\{j : 1 \leq j \leq L, j \neq i, C_j \text{ meets } R\}|$, counting indices throughout the capped exposure, including future indices. $V = \sum_{1 \leq j \leq L, j \neq i} \mathbf{1}_{\{|\text{vars}(C_j) \cap R| \geq 2\}}$ counts collision indices throughout the entire capped exposure and is $\mathcal{E}$ -measurable.

Conditional on $\mathcal{E}$, the hidden ordinary clauses are independent uniform proper 3-clauses on the nonroots, and the hidden nonroot parts of the test clauses are independent uniform proper 2-clauses there. These two collections are also independent of each other. To see this, condition first on the root subset and its signs separately at each index. Every permitted signed nonroot part remains equiprobable, and the original product law across indices is preserved. Revealing the remaining contents at the other indices does not change the law of the hidden coordinates.

For $i \leq m \leq L$, define

$$\begin{aligned} \mathcal{G}_m &= \mathcal{E} \vee \sigma(\text{ordinary contents through index } m), \\ \mathcal{H}_m &= \mathcal{G}_m \vee \sigma(\text{test parts through index } m). \end{aligned}$$

The set S_m is $\mathcal{G}_m$ -measurable: all test clauses are already true under w , whatever their nonroot parts. The event A_m is $\mathcal{H}_m$ -measurable, because this larger sigma-field knows all clauses of B_m . Future ordinary contents remain fresh given $\mathcal{H}_m$. Finally, put

$$p_m = q_d(S_m). \quad (3.4)$$

The quantity p_m is $\mathcal{G}_m$ -measurable and nondecreasing in m . Keeping the tests out of $\mathcal{G}_m$ will allow us to estimate pivotality after using $\mathcal{H}_m$ to estimate the remaining duration.

Remark (notation). $\mathcal{G}_m = \mathcal{E} \vee \sigma(\text{ordinary contents through index } m)$ is defined for integer indices $i \leq m \leq L$; the hidden nonroot parts of test clauses are not revealed in this filtration. $\mathcal{H}_m = \mathcal{G}_m \vee \sigma(\text{test parts through index } m)$ is defined for integer indices $i \leq m \leq L$ and includes the complete future clause-type information already contained in $\mathcal{E}$. For $i \leq m \leq L$, $p_m = q_d(S_m)$ is $\mathcal{G}_m$ -measurable and nondecreasing in m .

3.2 Necessary tests for root flips

Lemma 3.2 (Root-test bound). *For $i \leq m < L$, conditional on $\mathcal{G}_m$,*

$$\mathbb{P}(A_m \mid \mathcal{G}_m) \leq \mathbf{1}_{\{S_m \neq \emptyset\}} p_m^3 \quad \text{on } \{V = 0\}, \quad (3.5)$$

$$\mathbb{P}(A_m \mid \mathcal{G}_m) \leq \mathbf{1}_{\{S_m \neq \emptyset\}} p_m^2 \quad \text{on } \{V = 1\}. \quad (3.6)$$

Proof. Call a root x *available at time m* if no collision through index m , excluding i , has its literal on x as the unique root literal true under w . A collision excludes at most one root, so at least $3 - V$ roots are available. Availability is $\mathcal{E}$ -measurable.

For an available root x , the nonroot parts of its test clauses through m must kill S_m on A_m . Otherwise, take $z \in S_m$ satisfying all these parts and flip only x in the assignment (w, z) . Every clause of B_m remains satisfied. A clause with no true root literal under w already has a satisfied nonroot literal, which is unchanged. A clause with a true root literal other than the one on x keeps that literal. Finally, a clause whose unique true root literal is on x cannot be a collision, by availability; it is therefore a test clause for x , and z satisfies its nonroot part. The flipped assignment also satisfies C_i , contradicting A_m .

Given $\mathcal{G}_m$, the test families for distinct roots remain independent collections of random proper 2-clauses, independent of S_m . Each family has at most $Z \leq d$ members, so its probability of killing S_m is at most $q_d(S_m) = p_m$. On $V = 0$ all three roots are available; on $V = 1$ at least two are. Taking these necessary independent events proves the two inequalities, including the cases $p_m = 0$ or $S_m = \emptyset$. $\square$

4 Occupation and squared deletion delays

We continue with i and C_i fixed. Our goal is $\mathbb{E}D_i^2 = O(1)$, uniformly in the value of the fixed clause. The root tests bound the chance that a delay is present; the potential will bound how long that delay can continue.

Lemma 4.1 (Conditional occupation estimate). *With $K = 7$, for every $i \leq m < L$,*

$$\mathbb{E} \left[\sum_{s=m}^{L-1} \mathbf{1}_{\{S_s \neq \emptyset\}} p_s^\beta \mid \mathcal{H}_m \right] \leq K d^\beta. \quad (4.1)$$

Remark (notation). s is an integer prefix index summed over $m \leq s < L$, where $i \leq m < L$.

Proof. Suppose $s + 1$ is ordinary, a fact known under $\mathcal{E}$. Then $S_{s+1} = S_s[C_{s+1}]$, and C_{s+1} is still a uniform proper 3-clause on the nonroots given $\mathcal{H}_s$. Lemma 2.2 gives

$$\mathbf{1}_{\{S_s \neq \emptyset\}} p_s^\beta \leq 3d^{\beta-1} \mathbb{E}[P_d(S_{s+1}) - P_d(S_s) \mid \mathcal{H}_s] + d^\beta u^{-3}.$$

Condition on $\mathcal{H}_m$ and sum over such s . Every potential increment is nonnegative, even at nonordinary steps. Thus, pathwise,

$$\sum_{\substack{m \leq s < L \\ s+1 \text{ ordinary}}} (P_d(S_{s+1}) - P_d(S_s)) \leq P_d(S_L) - P_d(S_m) \leq d.$$

The tower property therefore bounds the contribution of the ordinary steps by $3d^\beta + Ld^\beta u^{-3}$. There are at most Z remaining steps: all indices $s + 1$ in question exceed i , and every nonordinary index meets a root. Each remaining summand is at most one. Since $Z \leq d \leq d^\beta$, $u = n - 3 \geq n/2$, and $n \geq 6$, the full sum is bounded by

$$3d^\beta + Ld^\beta u^{-3} + Z \leq \left(4 + \frac{80}{n^2}\right) d^\beta \leq 7d^\beta.$$

□

Corollary 4.2 (Remaining duration). *For $i \leq m < L$,*

$$\mathbb{E} \left[\sum_{s=m}^{L-1} \mathbf{1}_{A_s} \mid \mathcal{H}_m \right] \leq W_m, \quad W_m = \begin{cases} \min\{L, Kd^\beta p_m^{-\beta}\}, & p_m > 0, \\ L, & p_m = 0. \end{cases} \quad (4.2)$$

In particular, W_m is measurable with respect to the smaller sigma-field $\mathcal{G}_m$.

Remark (notation). For $i \leq m < L$, W_m is the $\mathcal{G}_m$ -measurable quantity $\min\{L, Kd^\beta p_m^{-\beta}\}$ when $p_m > 0$, and L when $p_m = 0$.

Proof. For $s \geq m$, we have $p_s \geq p_m$, and A_s implies $S_s \neq \emptyset$. If $p_m > 0$, divide (4.1) by p_m^β . The count is always at most L , giving both the minimum and the case $p_m = 0$. Measurability follows from that of d and p_m . □

The two estimates now have the compatible conditioning we need: the duration estimate knows whether A_m has occurred, while its upper bound W_m does not reveal the test parts. We can therefore multiply it by the independent-test probability from Lemma 3.2.

Proposition 4.3 (Squared delay conditional on the exposure). *For every $1 \leq i \leq L$,*

$$\mathbb{E}[D_i^2 \mid \mathcal{E}] \leq 2K^2 d^3 \mathbf{1}_{\{V=0\}} + 2KLd^{3/2} \mathbf{1}_{\{V=1\}} + L^2 \mathbf{1}_{\{V \geq 2\}}. \quad (4.3)$$

Proof. For $i = L$ we have $D_i = 0$. Otherwise, (3.3) implies

$$D_i^2 \leq 2 \sum_{m=i}^{L-1} \mathbf{1}_{A_m} \sum_{s=m}^{L-1} \mathbf{1}_{A_s}.$$

First condition on $\mathcal{H}_m$, where A_m is known, and apply Corollary 4.2. Then condition on $\mathcal{G}_m$, where W_m is known. Conditioning both estimates down to $\mathcal{E}$ gives

$$\mathbb{E}[D_i^2 \mid \mathcal{E}] \leq 2 \sum_{m=i}^{L-1} \mathbb{E}[W_m \mathbb{P}(A_m \mid \mathcal{G}_m) \mid \mathcal{E}]. \quad (4.4)$$

On $V = 0$, Lemma 3.2 bounds the product inside by

$$Kd^\beta \mathbf{1}_{\{S_m \neq \emptyset\}} p_m^{3-\beta} = Kd^\beta \mathbf{1}_{\{S_m \neq \emptyset\}} p_m^\beta.$$

This remains valid for $p_m = 0$, because the conditional probability then vanishes. Lemma 4.1 with $m = i$, conditioned further down to $\mathcal{E}$, bounds the sum in (4.4) by $2K^2 d^{2\beta} = 2K^2 d^3$. On $V = 1$, the same product is at most $Kd^\beta p_m^{2-\beta} \leq Kd^\beta$ when $p_m > 0$, and is zero when $p_m = 0$. There are at most L terms. Finally, on $V \geq 2$ use $D_i \leq L$. The cases are $\mathcal{E}$ -measurable, proving (4.3). $\square$

4.1 Averaging the collision costs

It remains to average the three terms of (4.3). Although one collision costs a factor L and two collisions cost L^2 , their probabilities are correspondingly small.

Lemma 4.4 (Root-incidence moments). *Uniformly in $n \geq 6$, i , and the fixed value of C_i ,*

$$\mathbb{E}d^3 = O(1), \quad \mathbb{E}[d^{3/2} \mathbf{1}_{\{V \geq 1\}}] = O(n^{-1}), \quad \mathbb{P}(V \geq 2) = O(n^{-2}). \quad (4.5)$$

Proof. Each of the other $L - 1$ independent clauses meets a specified root with probability $3/n$, and contains a specified pair of roots with probability $6/[n(n-1)]$. Union bounds over the three roots and their three pairs give incidence and collision probabilities satisfying

$$p_{\text{hit}} \leq \frac{9}{n}, \quad p_{\text{coll}} \leq \frac{18}{n(n-1)}.$$

The incidence count Z is binomial, so

$$\mathbb{E}e^Z = (1 + (e-1)p_{\text{hit}})^{L-1} \leq \exp(90(e-1)).$$

This uniformly bounded exponential moment gives $\mathbb{E}d^3 = O(1)$.

For each $j \neq i$ with $j \leq L$, let J_j be the event that C_j is a collision. Bound $\mathbf{1}_{\{V \geq 1\}}$ by $\sum_{j \neq i} \mathbf{1}_{J_j}$. Conditional on J_j , the count Z is one plus a binomial variable with $L - 2$ trials and success probability p_{hit} ; all other clauses retain their original law. Its exponential moment, and hence $\mathbb{E}[d^{3/2} \mid J_j]$, is bounded uniformly. Therefore

$$\mathbb{E}[d^{3/2} \mathbf{1}_{\{V \geq 1\}}] \leq \sum_{j \neq i} \mathbb{P}(J_j) \mathbb{E}[d^{3/2} \mid J_j] \leq O(1)(L-1)p_{\text{coll}} = O(n^{-1}).$$

Finally, independence at different indices gives

$$\mathbb{P}(V \geq 2) \leq \mathbb{E} \binom{V}{2} = \binom{L-1}{2} p_{\text{coll}}^2 = O(n^{-2}).$$

$\square$

Remark (notation). For $1 \leq j \leq L$ with $j \neq i$, J_j is the event that C_j meets at least two variables of R . p_{coll} is the common probability $\mathbb{P}(J_j)$ for any $1 \leq j \leq L$ with $j \neq i$; it depends only on n .

Combining Proposition 4.3 and Lemma 4.4, and using $L = 10n$, proves $\mathbb{E}D_i^2 = O(1)$ in the law conditional on C_i , with a constant independent of its value. We may now remove that conditioning. Lemma 3.1 gives

$$\text{Var}(\min\{H_n, 10n\}) = O(n) \quad (n \geq 6). \quad (4.6)$$

This completes the deletion argument. No information about the location of the satisfiability threshold was needed.

5 The uncapped hitting time and the lower bound

5.1 Removing the cap

For a fixed assignment, each independent clause is satisfied with probability $7/8$. A union bound over the 2^n assignments gives, for every $n \geq 3$ and integer $m \geq 0$,

$$\mathbb{P}(H_n > m) = \mathbb{P}(F_m \text{ is satisfiable}) \leq 2^n (7/8)^m. \quad (5.1)$$

In particular H_n is finite almost surely and has a finite second moment. For $n \geq 6$, write $T = \min\{H_n, L\}$ as before. The integer-valued tail-sum formula yields

$$\begin{aligned} \mathbb{E}(H_n - T)^2 &= \sum_{j=0}^{\infty} (2j+1) \mathbb{P}(H_n > L+j) \\ &\leq 120 \left(2(7/8)^{10}\right)^n. \end{aligned} \quad (5.2)$$

Here $\sum_{j \geq 0} (2j+1)r^j = (1+r)/(1-r)^2 = 120$ for $r = 7/8$, and $2(7/8)^{10} < 1$. Since

$$\text{Var}(H_n) \leq 2 \text{Var}(T) + 2 \text{Var}(H_n - T) \leq 2 \text{Var}(T) + 2\mathbb{E}(H_n - T)^2,$$

Equations (4.6) and (5.2) give the upper bound in Theorem 1.1 for $n \geq 6$. For $3 \leq n < 6$, the same tail bound gives $\mathbb{E}H_n^2 \leq 120 \cdot 2^n$, so enlarging the absolute constant covers these finitely many sizes.

5.2 Wilson's separation theorem

Define the survival function and its integer level locations by

$$s_n(m) = \mathbb{P}(H_n > m), \quad r_n(p) = \min\{m \geq 0 : s_n(m) \leq p\} \quad (0 < p < 1).$$

These locations are finite by (5.1) and positive because $s_n(0) = 1$. We use the following consequence of Wilson's theorem for the independent proper-clause model with replacement.

Remark (notation). For each $n \geq 3$ and $0 < p < 1$, $r_n(p) = \min\{m \geq 0 : s_n(m) \leq p\}$ is a finite positive integer determined by the unconditional distribution of H_n .

Theorem 5.1 (Wilson [7], Corollary 4). *Fix $1 > p_1 > p_2 > 0$. There is a constant $a(p_1, p_2) > 0$ such that, for all sufficiently large n , if integers m_1, m_2 satisfy $s_n(m_1) \geq p_1$ and $s_n(m_2) \leq p_2$, then*

$$m_2 - m_1 \geq a(p_1, p_2) \sqrt{n}.$$

Remark (notation). p_1 is a fixed deterministic probability level, independent of n , with $0 < p_2 < p_1 < 1$. p_2 is a fixed deterministic probability level, independent of n , with $0 < p_2 < p_1 < 1$. m_1 is a deterministic nonnegative integer endpoint satisfying $s_n(m_1) \geq p_1$; it may depend on n and need not be the integer level location $r_n(p_1)$. m_2 is a deterministic nonnegative integer endpoint satisfying $s_n(m_2) \leq p_2$; it may depend on n and need not be the integer level location $r_n(p_2)$. $a(p_1, p_2) > 0$ is a deterministic constant depending only on the fixed levels p_1, p_2 , independent of n and of the admissible endpoints m_1, m_2 .

Take $a_n = r_n(3/4)$ and $b_n = r_n(1/4)$. Minimality gives $s_n(a_n - 1) > 3/4$, whereas $s_n(b_n) \leq 1/4$. Theorem 5.1, applied at $a_n - 1$ and b_n , therefore implies $b_n - a_n + 1 \geq a(3/4, 1/4) \sqrt{n}$. Thus

$$b_n - a_n \geq c_0 \sqrt{n} \quad (5.3)$$

for an absolute $c_0 > 0$ and all sufficiently large n . Moreover,

$$\mathbb{P}(H_n \leq a_n) \geq \frac{1}{4}, \quad \mathbb{P}(H_n \geq b_n) = s_n(b_n - 1) > \frac{1}{4}.$$

Let H'_n be an independent copy of H_n . The two events $\{H_n \leq a_n, H'_n \geq b_n\}$ and $\{H'_n \leq a_n, H_n \geq b_n\}$ are disjoint for large n , each has probability at least $1/16$, and on each $|H_n - H'_n| \geq b_n - a_n$. Consequently,

$$\text{Var}(H_n) = \frac{1}{2} \mathbb{E}(H_n - H'_n)^2 \geq \frac{(b_n - a_n)^2}{16} \geq \frac{c_0^2}{16} n.$$

This completes the proof of Theorem 1.1.

Remark (notation). $c_0 > 0$ is an absolute constant, independent of n , such that $b_n - a_n \geq c_0 \sqrt{n}$ for all sufficiently large n . H'_n has the same distribution as H_n and is independent of H_n .

Corollary 5.2 (Fixed central windows). *For every fixed $0 < \eta < 1/2$,*

$$r_n(\eta) - r_n(1 - \eta) = \Theta_\eta(\sqrt{n}).$$

Proof. The lower bound follows from Theorem 5.1 at $r_n(1 - \eta) - 1$ and $r_n(\eta)$, absorbing the additive one. For the upper bound, put $\mu_n = \mathbb{E}H_n$ and choose $t = (2Cn/\eta)^{1/2}$, with C from Theorem 1.1. Chebyshev's inequality gives $\mathbb{P}(|H_n - \mu_n| \geq t) \leq \eta/2$. For every nonnegative integer $m \leq \mu_n - t$ this implies $s_n(m) \geq 1 - \eta/2 > 1 - \eta$, whereas at $m = \lceil \mu_n + t \rceil$ it implies $s_n(m) \leq \eta/2 < \eta$. Hence $r_n(1 - \eta) > \mu_n - t$ and $r_n(\eta) \leq \lceil \mu_n + t \rceil$, proving the upper bound. If $\mu_n - t < 0$, the first inequality follows instead from $r_n(1 - \eta) \geq 0$. $\square$

The variance bound controls fluctuations around the finite-size mean $\mathbb{E}H_n$. Neither the proof nor Corollary 5.2 requires a limiting value of $\mathbb{E}H_n/n$ or a limiting distribution for the centered hitting time.

References

- [1] S. Boucheron, O. Bousquet, G. Lugosi, and P. Massart. Moment inequalities for functions of independent random variables. *The Annals of Probability*, 33(2):514–560, 2005.
- [2] G. Carenini. A polynomial scaling window for random k -SAT and a proof of the satisfiability conjecture. Technical Report TR26-229, Electronic Colloquium on Computational Complexity, 2026. Published 5 October 2026. <https://eccc.weizmann.ac.il/report/2026/229/>.
- [3] G. Carenini. The scaling window of random k -SAT. Technical Report TR26-145, Electronic Colloquium on Computational Complexity, 2026.
- [4] B. Efron and C. Stein. The jackknife estimate of variance. *The Annals of Statistics*, 9(3):586–596, 1981.
- [5] E. Friedgut. Sharp thresholds of graph properties, and the k -SAT problem. *Journal of the American Mathematical Society*, 12(4):1017–1054, 1999. With an appendix by Jean Bourgain.
- [6] OpenAI. Variance of the Random k -SAT Hitting Time. OpenAI Math Release preprint [OAI:Variance-of-the-Random-k-SAT-Hitting-Time-September-27-2026](#), 2026.
- [7] D. B. Wilson. On the critical exponents of random k -SAT. *Random Structures & Algorithms*, 21(2):182–195, 2002.