Definability on a Random 3-CNF Formula

Albert Atserias*
Universitat Politècnica de Catalunya
atserias@lsi.upc.edu

Abstract

We consider the question of certifying unsatisfiability of random 3-CNF formulas. At which densities can we hope for a simple sufficient condition for unsatisfiability that holds almost surely? We study this question from the point of view of definability theory. The main result is that first-order logic cannot express any sufficient condition that holds almost surely on random 3-CNF formulas with $n^{2-\alpha}$ clauses, for any irrational positive number α. In contrast, it can when the number of clauses is $n^{2+\alpha}$, for any positive α. As an intermediate step, our proof exploits the planted distribution for 3-CNF formulas in a new technical way. Moreover, the proof requires us to extend the methods of Shelah and Spencer for proving the zero-one law for sparse random graphs to arbitrary relational languages.

1. Introduction

The complexity of 3-SAT on random instances has received a good deal of attention in recent years in many different areas such as satisfiability testing [28], propositional proof complexity [11], statistical physics methods applied to combinatorial optimization [27], integer and linear programming [10], and hardness of approximation [18]. Perhaps the main question is whether there exists a simple (polynomial-time) property of 3-CNF formulas that guarantees unsatisfiability and holds for typical unsatisfiable formulas [6, 21, 19, 18]. The positive answer would provide a good deal of information about the structure of unsatisfiable formulas. The negative answer would imply a strong hardness result for 3-SAT. Following [18], this scene of hardness is called the random 3-SAT hypothesis.

We work in the random model of 3-CNF formulas in which each possible clause is chosen with independent probability p, so the expected number of clauses is $m = \Theta(n^3 p)$, where n is the number of variables. The choice of

this model, instead of the so-called model B in which the number of clauses is fixed, is only a matter of convenience and does not affect the significance of the results. Clearly, whether the random 3-SAT hypothesis holds or not depends on m. If m is very large, say $m = \Omega(n^3)$, then it is almost trivial that a random 3-CNF formula is almost surely (a.s.) unsatisfiable for the simple reason that it contains a constant-size unsatisfiable subformula. On the contrary, if m is very small, say $m = O(1)$, then it is not hard to see that a random 3-CNF formula is a.s. satisfiable, so a sufficient condition for unsatisfiability cannot hold a.s. The cases of most interest occur when m is close to the point where a random 3-CNF formula undergoes the transition from being a.s. satisfiable to being a.s. unsatisfiable. Theoretical results establish that the transition occurs somewhere between $3.42n$ and $4.51n$ [22, 12], and experimental results suggest that it occurs around $m = 4.2n$ [28].

A motivating example Before we describe our results, let us start with a motivating example from the theory of random graphs $G(n, p)$. Consider planarity. It is known that if $p = O(n^{-1-\epsilon})$, then a random graph is a.s. planar, and if $p = \Omega(n^{-1+\epsilon})$, then it is a.s. non-planar. In fact, the threshold is much finer, but let us ignore this for the moment. How can we certify the non-planarity of a typical graph when $p = \Omega(n^{-1+\epsilon})$? By the easy direction of Kuratowski's Theorem, it suffices to find a subgraph that is homeomorphic to $\mathbf{K}_5$ or $\mathbf{K}_{3,3}$. Let us argue that it suffices to consider subgraphs of bounded size. Indeed, consider the graph $\mathbf{K}_{3,3}^\epsilon$ that results from a $\mathbf{K}_{3,3}$ by subdividing its nine edges into nine disjoint paths of length $\lceil 1/3\epsilon \rceil$. The density of any subgraph of $\mathbf{K}_{3,3}^\epsilon$ is strictly below $\frac{1}{1-\epsilon}$, so it follows from the seminal work of Erdös and Renyi [16] that $\mathbf{K}_{3,3}^\epsilon$ occurs as a subgraph in $G(n, p)$ a.s. when $p = \Omega(n^{-1+\epsilon})$. Clearly $\mathbf{K}_{3,3}^\epsilon$ is homeomorphic to $\mathbf{K}_{3,3}$, and for fixed ϵ, the size of $\mathbf{K}_{3,3}^\epsilon$ is also fixed. So the non-planarity of $G(n, p)$ is certified a.s. by the existence of a fixed-size subgraph that is homeomorphic to $\mathbf{K}_{3,3}$.

Testing for the existence of a fixed-size subgraph is clearly efficient, certainly in polynomial-time. As it turns out, the existence of fixed-size substructures is the paradig-

matical example of a property that is definable in first-order logic, which of course is much stronger in definability power, but still polynomial-time. In turn, first-order definability on random structures is a well-established topic with many deep results (see [32]). In view of the situation with planarity on random graphs, it is thus scholarly necessary to consider first-order definability on random 3-CNF formulas, with a focus on certificates for unsatisfiability.

Results of this paper   3-CNF formulas are viewed as structures $M = (V, R_0, \ldots, R_7)$ whose domain V is the set of propositional variables, and whose relations $R_i \subseteq V^3$ define which clauses, of each of the eight possible types, appear in the formula. Here $n = |V|$ and $m = \sum_i |R_i|$. At which densities is there a sufficient condition for unsatisfiability that is first-order definable and holds a.s.? For such densities we say that first-order logic *can certify* unsatisfiability a.s. When $m = \Omega(n^{2+\alpha})$, with $\alpha > 0$, it is not hard to see that a random 3-CNF formula with m clauses contains a constant-size unsatisfiable subformula a.s. Clearly, a first-order sentence can simply express this, so first-order logic can certify unsatisfiability a.s. at this density. Our main result establishes that this is essentially optimal:

Theorem 1 *Let F be a random 3-CNF formula with n variables and m expected clauses.*

1. *If $m = \Theta(n^{2+\alpha})$ with arbitrary $\alpha > 0$, then first-order logic can certify unsatisfiability a.s.*

2. *If $m = \Theta(n^{2-\alpha})$ with irrational $\alpha > 0$, then first-order logic cannot certify unsatisfiability a.s.*

Requiring α to be irrational in 2. is a technicality that will be discussed later in the paper. The main conceptual contribution of this work is in studying first-order definability on a random 3-CNF with a focus on certificates for unsatisfiability, where we succeed in establishing a tight result. The main technical contribution is in the the proof of the non-expressibility result. The main new idea is a new approach for finding finite structures that are first-order indistinguishable. To our knowledge, this approach is novel. We consider the so-called *planted distribution* for random 3-CNF formulas and argue that they are indistinguishable from regular random 3-CNF formulas a.s. We note that the planted distribution appears only as a technical detour in the proof; the result itself is about the usual distribution on 3-CNF formulas. This aspect of the proof is interesting; let us discuss it in more detail.

Proof techniques   Proving the inexpressibility result required us to study definability on random 3-CNF formulas. The approach we follow is related to that taken by others for studying first-order logic on random graphs (we discuss related work in the next subsection). However, establishing the main result required some new techniques. The first idea that may come to mind is the following: if we take a random 3-CNF formula at a density below the threshold of satisfiability, say at $m = o(n)$, and one at a density above it, say at $m = \omega(n)$, then a.s. one is satisfiable and the other is not, and with some hope perhaps, a first-order sentence may be unable to distiguish between them. Unfortunately, this argument does not work. The reason it does not is that two random 3-CNF formulas with a number of clauses asymptotically below n, and above n respectively, have different occurrences of constant size subformulas. Therefore, a first-order sentence can always be designed to distinguish between them. This suggests that we take both random formulas with the same density, which in turn, spoils the property that one is satisfiable and the other is not.

The new approach that works is the following. We want to pick our formulas in such a way that a.s. one is satisfiable and the other is not, and, among other things, both have the same constant-size subformulas. Here is one choice that guarantees this: take one formula from the usual distribution and the other from the planted distribution at the *same* density. The planted distribution $P(n, p)$ is the one in which each clause that is not falsified by a fixed *planted truth assignment* is chosen with probability p, and the rest of clauses are banned. The planted distribution guarantees that the resulting 3-CNF formula is satisfiable, and showing that it has the same constant-size subformulas as a formula from the usual distribution is one of the main steps of our proof. In fact, we need to show much more: the two formulas are a.s. indistinguishable by first-order formulas of any fixed quantifier depth. The conclusion is that a first-order property cannot certify unsatisfiability a.s.

Related work   Random 3-CNF formulas were introduced a long time ago. See [13] for surveys. The results of Chvátal and Szemerédi [11] were the first to indicate a certain degree of typical-case hardness. These were extended by subsequent work [6, 8, 7, 4]. Despite these early results, the first to explicitly formulate the random 3-SAT hypothesis was Feige [18], who established that, assuming the hypothesis is true at $m = O(n)$, one could take it as a hardness assumption for showing a number of non-approximability results that do not seem to follow from PCP constructions. Using spectral techniques, Friedman and Goerdt [19] found a polynomial-time certifying algorithm at $m = \Omega(n^{3/2+\epsilon})$. Before [19], Beame, Karp, Pitassi and Saks [6] already had a polynomial-time certifying algorithm at $m = \Omega(n^2/\log n)$. Extending the techniques in [19] is a matter of ongoing research.

First-order definability on random structures is a well-studied topic. The pioneers were Glebskii et al. [20] and Fagin [17] with the 0-1 law. Their work was followed by

[26, 24, 29, 30] and many others. The techniques for dealing with sparse structures such as $G(n, n^{-\alpha})$ were introduced by Shelah and Spencer and influenced much of the subsequent work, including this. Sufficient conditions for hard properties of graphs were also studied in [9]. The planted distribution for 3-CNF formulas was considered in [2] with a completely different goal.

The current work requires comparison to our previous work [3]. In [3] we showed that Datalog cannot certify unsatisfiability a.s. at $m = O(n^{3/2-\epsilon})$. Datalog and first-order logic are rather orthogonal. While both logics define polynomial-time properties only, Datalog cannot define non-monotone properties and first-order logic cannot define non-local properties (see [14]). The non-expressibility results we obtain are thus incomparable. The results about Datalog indicate that *unbounded recursive existential quantification* is unable to certify unsatisfiability, and the results about first-order logic indicate that so is *bounded quantifier alternation*. While neither result can really be seen as strong evidence towards the random 3-SAT hypothesis, the Datalog result generalizes the lower bound for resolution in [11], but the first-order logic result does not seem to have a proof complexity counterpart yet. The techniques for proving both results are also totally different.

2. Preliminaries

Relational languages and structures A relational language $L = \{R_1, R_2, \ldots\}$ is a set of relation symbols each with an associated arity. An L-structure is a tuple $M = (U, R_1^M, R_2^M, \ldots)$ where U is the universe, and R_i^M is a relation over U of the arity of the symbol R_i. We identify structures and their universe when this does not lead to confusion. First-order formulas are formed from atomic formulas of the form $R_i(x_1, \ldots, x_{r_i})$ and $x_i = x_j$ by means of negations, conjunctions, disjunctions, and existential and universal quantification over first-order variables. Here, the variables x_i range over the elements of the universe. The semantics of first-order logic can be found in any standard textbook in logic such as [15].

Let $L_8 = \{R_0, \ldots, R_7\}$ be the relational language of eight relation symbols of arity three. Observe that a 3-CNF formula is nothing but an L_8-structure: let its universe be the set of propositional variables, and let the tuples of its relations indicate which 3-clauses of each of the eight different types appear in the formula. For example, if the clause $v_1 \lor v_2 \lor v_3$ is in the formula, add (v_1, v_2, v_3) to R_0, if the clause $\neg v_1 \lor \neg v_2 \lor \neg v_3$ is in the formula, add (v_1, v_2, v_3) to R_7, etc. Conversely, structures for L_8 are nothing but 3-CNF formulas by reversing the interpretation.

Formulas and their Probability Spaces As in the random graph model of Erdös and Renyi, there are two main families of distributions of interest. For a fixed number of variables n, the first family considers the number of clauses m as fixed, and endows the space of k-CNF formulas with m clauses on n variables with the uniform distribution. The second family considers each clause on n variables independently with probability p, and endows the space of all k-CNF formulas on n variables with the product distribution. Both these families have several variants according to whether clauses are ordered tuples or sets, and may, or may not, have repeated or complementary literals. As in the random graph model again, which space to use is often a matter of convenience, and rarely an important issue as far as the results are concerned.

Since we have adopted the framework in which k-CNF formulas are structures over a relational language, it is convenient to define a probability space on finite structures. For k-CNF formulas, the distribution we choose turns out to be the product distribution in which clauses are ordered tuples possibly with repeated and complementary literals.

Definition 1 *Let $L = \{R_1, \ldots, R_r\}$ be a relational language, let r_i be the arity of R_i, and let $p = p(n)$ be such that $0 \le p \le 1$. Let $L(n)$ be the class of all L-structures with universe $\{1, \ldots, n\}$, and let $L(n, p)$ be the probability distribution on $L(n)$ that assigns probability*

$$\prod_{i=1}^{r} p^{|R_i^M|}(1-p)^{|U^{r_i}|-|R_i^M|}$$

to each $M = (U, R_1^M, \ldots, R_r^M)$ in $L(n)$.

For the proof of the main technical result of the paper we will need a detour through the *planted distribution* for 3-CNF formulas. In a nutshell, the planted distribution consists in drawing 3-CNF formulas in the usual way, except that the clauses that falsify a fixed *planted* truth assignment are forbidden (have zero probability). For concreteness, the planted truth assignment is always the same, namely, the one that assigns "false" to the first half of the variables and "true" to the rest. Formally,

Definition 2 *Let $p = p(n)$ be such that $0 \le p \le 1$. Let f be the truth assignment on the variables $x_1, \ldots, x_n$ that assigns false to the first $\lfloor n/2 \rfloor$ variables and true to the rest. Let $P_8(n, p)$ be the probability distribution on $L_8(n)$ that is obtained by adding each clause that is not falsified by f independently, with probability p.*

Obviously, the set of 3-CNF formulas with non-zero probability in $P_8(n, p)$ are exactly those satisfied by the planted truth assignment. It is not hard to see that the number of clauses that do not falsify the planted truth assignment is $7n^3$.

3. Counting extensions: multivariate polynomial method

A hypergraph $H = (V, E)$ consists of a set of vertices V and a set E of sets of vertices called hyperedges. A weighted hypergraph is a hypergraph (V, E) together with a weight function $w : E \to \mathbf{R}$ that assigns a real number $w(e)$ to each hyperedge $e \in E$. Let $H = (V, E)$ be a hypergraph with positive weights on its hyperedges, and let $n = |V|$. Consider a collection $\{X_u : u \in V\}$ of independent random variables, where each X_u is either a $\{0, 1\}$-random variable with expected value p_u, or the constant random variable $X_u = p_u$. Here $0 \le p_u \le 1$ for every $u \in V$. Let Y be the following polynomial:

$$Y = \sum_{e \in E} w(e) \prod_{u \in e} X_u.$$

For every $A \subseteq V$, let Y_A be the partial derivative of Y with respect to $\{X_u : u \in A\}$. For every $i \ge 0$, let $E_i = \max\{E(Y_A) : A \subseteq V, |A| = i\}$, where $E(Y_A)$ denotes expectation.

Theorem 2 (Corollary 4.1.3 in [25]) *If there is a constant $\gamma > 0$ such that $E_i/E_0 = O(n^{-\gamma})$ for all $i > 0$, then there are constants $\epsilon > 0$ and $\delta > 0$ such that*

$$\Pr\Big[|Y - E(Y)| > n^{-\epsilon}E(Y)\Big] < e^{-n^\delta}.$$

A prototypical application of Theorem 2 is the estimation of the number of occurrences of a small subgraph G in a random graph on n nodes. In that case, V is the set of possible edges on $\{1, \ldots, n\}$, and E is the set of possible placements of G in $\{1, \ldots, n\}$. In fact, the rest of the argument is a particular case of Theorem 3 below, so we turn to proving that immediately.

We extend the notion of rooted graph from [29, 30] to arbitrary languages. Let us fix a finite relational language L. A *rooted structure*, or *extension*, is a structure H with a designated subset R of elements, called the *roots*. We denote it by (R, H). The *type* of a rooted structure is (v, e), where v is the number of elements of H that are not roots, and e is the number of tuples in the relations of H with at least one non-root element. Its density is e/v. The set of root elements is usually denoted by the ordered tuple $(P_1, \ldots, P_r)$, and the set of non-root elements is denoted by the ordered tuple $(Q_1, \ldots, Q_v)$. There will always be such an implicit ordering on the elements of H that will be clear from context. Let M be any structure, let $\mathbf{x} = (x_1, \ldots, x_r)$ and $\mathbf{y} = (y_1, \ldots, y_v)$ be tuples of distinct elements of M. We let $g(\mathbf{x}, \mathbf{y})$ be the mapping $P_i \mapsto x_i$ and $Q_i \mapsto y_i$. We say that $\mathbf{y}$ is an (R, H)-extension of $\mathbf{x}$ in M if $g(\mathbf{x}, \mathbf{y})$ maps tuples in the relations of H with at least one non-root element to tuples of the same relation in M. If $\mathbf{t}$ is a tuple of elements of H, we let $\mathbf{t}(\mathbf{x}, \mathbf{y})$ be the image of $\mathbf{t}$ under $g(\mathbf{x}, \mathbf{y})$.

Recall that we identify structures with their universes when this does not lead to confusion. Let $R \subset S \subseteq H$. We call (R, S) a *subextension* of (R, H). Let $R \subseteq S \subset H$. We call (S, H) a *nailextension* of (R, H). Observe that a rooted structure is a subextension and a nailextension of itself.

The next result analyzes the number of extensions of the tuples of a random structure. For the case of undirected graphs, this is the hardest result in the proof of [32]. Its proof used Janson's inequality together with a number of ad-hoc arguments. An alternative transparent proof was obtained by Kim and Vu [23] using the multivariate polynomial method, still for undirected graphs. We extend the result for general structures, where the machinery of Kim and Vu is very useful.

Theorem 3 *Let (R, H) be a rooted structure of type (v, e), and let ρ be the maximal density of its subextensions. Let $\alpha > 0$, let $M \sim L(n, p)$ where $p = \Theta(n^{-\alpha})$, and for every tuple $\mathbf{x}$ of v distinct elements in M, let $N(\mathbf{x})$ be the number of (R, H)-extensions of $\mathbf{x}$ in M.*

(i) If $\alpha > 1/\rho$, then $N(\mathbf{x}) = 0$ for some $\mathbf{x}$, almost surely,

(ii) If $\alpha < 1/\rho$, then $N(\mathbf{x}) > 0$ for every $\mathbf{x}$, almost surely.

Moreover, $E(N(\mathbf{x})) = \Theta(n^{v - \alpha e})$, and in case (ii) $N(\mathbf{x}) \sim E(N(\mathbf{x}))$ for every $\mathbf{x}$, almost surely.

Proof: Fix an arbitrary tuple $\mathbf{x} = (x_1, \ldots, x_r)$ of $r = |R|$ distinct elements of M. For every subextension (R, S), let (v_S, e_S) be its type, and let $N_S(\mathbf{x})$ be the number of (R, S)-extensions of $\mathbf{x}$ in M. Note that $N(\mathbf{x}) = N_H(\mathbf{x})$ and that if $N_S(\mathbf{x}) = 0$, then $N_H(\mathbf{x}) = 0$. Also, $E(N_S(\mathbf{x})) = \Theta(n^{v_S}n^{-\alpha e_S})$. If $\alpha > 1/\rho$, then $E(N_S(\mathbf{x})) = o(1)$ for the subextension (R, S) of maximal density $\rho = e_S/v_S$. In this case, $N_S(\mathbf{x}) = 0$ almost surely by Markov's inequality, so $N_H(\mathbf{x}) = 0$ almost surely.

Suppose now $\alpha < 1/\rho$. We aim for an application of Theorem 2. For every relation symbol $R_j \in L$ of arity s and every s-tuple $\mathbf{t}$ of elements in M, let $Z(j, \mathbf{t})$ be the random variable indicating whether $\mathbf{t} \in R_j^M$ or not. Note that

$$N_S(\mathbf{x}) = \sum_{\mathbf{y}} \prod_{j, \mathbf{t}} Z(j, \mathbf{t}(\mathbf{x}, \mathbf{y})),$$

with $\mathbf{y} = (y_1, \ldots, y_{v_S})$ ranging over all tuples of v_S distinct elements of $M - \{x_1, \ldots, x_r\}$, and j and $\mathbf{t}$ ranging over all pairs such that $\mathbf{t} \in R_j^S$ and $\mathbf{t}$ has some non-root element. Recall that $\mathbf{t}(\mathbf{x}, \mathbf{y})$ denotes the image of $\mathbf{t}$ under the mapping $P_i \mapsto x_i$ and $Q_i \mapsto y_i$, where $R = \{P_1, \ldots, P_r\}$ and $S = \{Q_1, \ldots, Q_{v_S}\}$. Let us bound $E_i(N_H(\mathbf{x}))$ for $1 \le i \le e$. Let γ be the minimum $v_S - \alpha e_S$ over all subextensions (R, S). Note that $\gamma > 0$ because ρ is the maximal density of the subextensions and $\alpha < 1/\rho$. For every set A of exactly i tuples of H involving some non-root element, let S_A be the non-root elements appearing in A, and

let j_A be the cardinality of S_A. Observe that $(R, R \cup S_A)$ is a subextension of (R, H) of type (j_A, i). Let j be the minimum over all j_A and let (R, S) be the corresponding subextension. Clearly $j - \alpha i \geq \gamma$. Moreover,

$$E_i(N_H(\mathbf{x})) = O(n^{v-j} n^{-\alpha(e-i)})$$

because having fixed i tuples, at most $v - j$ nodes are left free. Note that $E_i / E_0 = O(n^{-\gamma})$. Therefore, Theorem 2 implies that

$$\Pr\left[|N_H(\mathbf{x}) - E(N_H(\mathbf{x}))| > n^{-\epsilon} E(N_H(\mathbf{x})) \right] < e^{-n^{\delta}}$$

for some constants $\epsilon > 0$ and $\delta > 0$. The conclusion follows because $e^{n^{\delta}}$ grows faster than any polynomial, and there are polynomially many possible $\mathbf{x}$. $\square$

When $R = \emptyset$, the rooted structure (R, H) can be identified with the structure H. Moreover, the number of $(\emptyset, H)$-extensions of the empty tuple $\mathbf{x} = \emptyset$ in M coincides with the number of copies of H in M. Thus, we have the following corollary.

Corollary 1 *Let H be a structure with v elements and e tuples, and let ρ be the maximal density of its substructures. Let $\alpha > 0$, let $M \sim L(n, p)$ where $p = \Theta(n^{-\alpha})$, and let N be the number of copies of H in M.*

(i) If $\alpha > 1/\rho$, then $N = 0$, almost surely,

(ii) If $\alpha < 1/\rho$, then $N > 0$, almost surely.

Moreover, $E(N) = \Theta(n^{v-\alpha e})$, and in case (ii) $N \sim E(N)$, almost surely.

Let us now concentrate on random 3-CNF formulas, and in particular, on the planted distribution $P_8(n, p)$. Our aim is to show that essentially the same result as in Theorem 3 holds for the planted distribution. The key difference is in the need to require $\alpha > 1$; this guarantees that the extension is satisfiable in case (ii) and makes the proof possible. Let us note that the proof requires an application of Tarsi's Lemma [1] and is significantly different for case (ii).

Theorem 4 *Let (R, H) be a rooted 3-CNF formula of type (v, e), and let ρ be the maximal density of its subextensions. Let $\alpha > 1$, let $M \sim P_8(n, p)$ where $p = \Theta(n^{-\alpha})$, and for every tuple $\mathbf{x}$ of v distinct elements in M, let $N(\mathbf{x})$ be the number of (R, H)-extensions of $\mathbf{x}$ in M.*

(i) If $\alpha > 1/\rho$, then $N(\mathbf{x}) = 0$ for some $\mathbf{x}$, almost surely,

(ii) If $\alpha < 1/\rho$, then $N(\mathbf{x}) > 0$ for every $\mathbf{x}$, almost surely.

Moreover, $E(N(\mathbf{x})) = \Theta(n^{v-\alpha e})$, and in case (ii) $N(\mathbf{x}) \sim E(N(\mathbf{x}))$ for every $\mathbf{x}$, almost surely.

Proof: For (i) we proceed as in Theorem 3. The expectation $E(N_S(\mathbf{x}))$ is again certainly bounded by $O(n^{v_S} n^{-\alpha e_S})$. Therefore, if $\alpha > 1/\rho$, then $N_S(\mathbf{x}) = 0$ almost surely for the S of maximal density, so $N_H(\mathbf{x}) = 0$ almost surely. The case (ii) requires a totally new proof.

Suppose $\alpha < 1/\rho$. Fix a tuple $\mathbf{x} = (x_1, \ldots, x_r)$ of distinct elements of M. Consider the (R, H)-extensions of $\mathbf{x}$. Note that some of the elements of $\mathbf{x}$ are in the first half of the variables and some are in the second. Since $\alpha > 1$, this implies that the maximum density of every subextension is strictly below 1. It follows from this that H is a satisfiable 3-CNF even when the variables in $\mathbf{x}$ are set to the truth value according to the half of the variables they belong to. Indeed, if it were unsatisfiable it would have a minimally unsatisfiable subformula whose density would be above 1 by Tarsi's Lemma [1], which contradicts the assumption $\alpha < 1/\rho$. We now use this fact in a crucial way. Fix a truth assignment to the non-root variables that satisfies H, and let a be the number of non-root variables of H that are set to false by this assignment, and let b be the number of non-root variables of H that are set to true by this assignment. Notice that $a + b = v_H$. Let us compute an upper bound and a lower bound for $E(N_H(\mathbf{x}))$. The number of possible occurrences of H is generously bounded by n^{v_H}. Hence $E(N_H(\mathbf{x})) = O(n^{v_H} p^{e_H})$. For the lower bound, note that the number of placements of H of non-zero probability is at least

$$\binom{n/2 - r}{a} \binom{n/2 - r}{b},$$

because any placement all whose clauses are satisfied by the planted truth assignment has non-zero probability. The probability of each such placement is at least $p^{e_H} (1 - p)^{e'_H}$ for some fixed constant e'_H that depends on H only. Hence, the expectation of $N_H(\mathbf{x})$ is at least

$$\binom{n/2 - r}{a} \binom{n/2 - r}{b} p^{e_H} (1 - p)^{e'_H}.$$

Since $a + b = v_H$, r is a constant, and $p = o(1)$, the upper and lower bounds are related by a constant, so $E(N_H(\mathbf{x})) = \Theta(n^{v_H} n^{-\alpha e_H})$.

The rest of the argument is now essentially the same argument as in Theorem 3. In this case, though, notice that some of the random variables $Z(j, \mathbf{t}(\mathbf{x}, \mathbf{y}))$ will be constants set to 0 because the corresponding clause will not be satisfied by the planted truth assignment. This is allowed in the hypothesis of Theorem 2. The computation of $E_i(N_H(\mathbf{x}))$ is the same since we only aim for an upper bound. The result follows. $\square$

4. More on Rooted Structures

The purpose of this section is to extend the concepts in [32] to general relational structures and the planted distribution. For this section we fix an irrational number $\alpha > 0$. Recall the definition of *rooted structure*, or *extension*, from Section 3.

Definition 3 *Let (R, H) be an extension of type (v, e). If $v - e\alpha$ is positive we call (R, H) sparse. If $v - e\alpha$ is negative we call (R, H) dense.*

Notice that since α is irrational, every extension is either sparse or dense; this will play a crucial role in the proofs. Recall that we identify structures with their universes when this does not lead to confusion. Let $R \subset S \subseteq H$. We call (R, S) a *subextension* of (R, H). Let $R \subseteq S \subset H$. We call (S, H) a *nailextension* of (R, H). Observe that a rooted structure is a subextension and a nailextension of itself.

Definition 4 *An extension is called rigid if all its nailextensions are dense. It is called safe if all its subextensions are sparse.*

Next we turn to the key concept of closure. A preliminary version of this concept was introduced in [29, 30] and used in [31]. Unfortunately, the original definition suffered some technical problems and required later refinement (see [32]). We extend it to general structures.

Definition 5 *Let M be a structure, let $\mathbf{x} = (x_1, \dots, x_i)$ be a tuple of elements of M, and let $t \geq 0$. The t-closure of $\mathbf{x}$, denoted $cl_t^M(\mathbf{x})$, is the smallest $X \subseteq M$ containing $\{x_1, \dots, x_i\}$ that does not have rigid extensions with at most t non-roots.*

The closure of $\mathbf{x}$ is constructively obtained by letting $X_0 = \{x_1, \dots, x_i\}$, and letting (X_j, X_{j+1}) be any rigid extension of X_j with at most t non-roots, if there is one. The last X_j is $cl_t^M(\mathbf{x})$. It is not hard to see that both approaches lead to the same set. Since the structure M is usually understood from context, we may drop the superscript in $cl_t^M(\mathbf{x})$. Here are a couple of properties of closures with easy proofs:

Lemma 1 *Let M be a structure, and let $\mathbf{x} = (x_1, \dots, x_i)$ be a tuple of distinct elements of M.*

(i) If $y \in cl_r(\mathbf{x}) - \mathbf{x}$, then $cl_t(\mathbf{x}, y) \subseteq cl_{r+t}(\mathbf{x})$.

(ii) If $|M| \leq t$ and $cl_t(\mathbf{x}) \neq M$, then $(cl_t(\mathbf{x}), M)$ is safe.

Proof: Statement (i) is clear. For (ii), if $(cl_t(\mathbf{x}), M)$ is not safe and $cl_t(\mathbf{x}) \neq M$, then it has a dense subextension $(cl_t(\mathbf{x}), S)$ since α is irrational. Let S be minimal with that property. Then $(cl_t(\mathbf{x}), S)$ is actually rigid for otherwise some nailextension (T, S) would be sparse and

$(cl_t(\mathbf{x}), T)$ would have to be dense contradicting the minimality of S. Here we used the irrationality of α again. But since $|S| \leq |M| \leq t$, necessarily $S \subseteq cl_t(\mathbf{x})$; a contradiction. □

The following lemma states that closures are almost surely bounded in size. Its proof is adapted from that of Theorem 6.2 in [32] for undirected graphs. We note that it also works for the planted distribution:

Lemma 2 *For every $i > 0$ and $t > 0$, there exists a constant $K > 0$ such that, for $M \sim L(n, p)$ or $M \sim P(n, p)$ with $p = \Theta(n^{-\alpha})$, almost surely $|cl_t(\mathbf{x})| < K$ for all tuples $\mathbf{x} = (x_1, \dots, x_i)$ of distinct elements of M.*

Proof: Let $\beta = \max\{(v - \alpha e)/v : v \leq t, \; v - \alpha e < 0\}$. Let $K > 0$ be such that $i + K\beta < 0$; since $\beta < 0$, such a K exists. The existence of a closure $cl_t(R)$ of size at least K for some $R \subseteq M$ of size i implies the existence of $R = S_0 \subseteq \cdots \subseteq S_l$ where each (S_j, S_{j+1}) is rigid with at most t non-roots and $K + r \leq S_l \leq K + r + t$. Let (v_j, e_j) be the type of (S_j, S_{j+1}). Since (S_j, S_{j+1}) is rigid, it is dense, so $v_j - \alpha e_j < 0$. Moreover, $\sum_j v_j \geq K$. Now, the expected number of such sequences of extensions is bounded by $O(n^{v - \alpha e})$ where $v = i + \sum_j v_j$ and $e = \sum_j e_j$. However,

$$
\begin{aligned}
v - \alpha e &= i + \sum_j v_j - \alpha \sum_j e_j = \\
&= i + \sum_j (v_j - \alpha e_j) \leq \\
&\leq i + \sum_j \beta v_j \leq \\
&\leq i + K\beta.
\end{aligned}
$$

Recall that $\beta < 0$ for the last inequality. Now, $i + K\beta < 0$ by our choice of K. Therefore, the number of such extensions is zero almost surely. Since there is a bounded number of possible sequences $S_0 \subseteq \cdots \subseteq S_l$ as above, the number of closures of size at least K is also zero almost surely. □

5. Proof of Main Result

In this section we prove Theorem 1. Since the first part of the theorem is much easier, let us concentrate on the second part first. This will require the argument that we sketched in the introduction and that we sketch again here. Draw two random 3-CNF formulas, one from the usual distribution and the other from the planted distribution which ensures that the formula is satisfiable. The key of the argument is that these two formulas are a.s. indistinguishable by first-order formulas of any fixed number of quantifers. The result will follow since then, a first-order formula is unable

to distinguish the first, which is almost surely unsatisfiable, from the second, which is always satisfiable.

Before we jump into the proof, we need some preparation. Two tuples with isomorphic t-closures can be viewed as having the same "special" extension properties. The term "special" is justified by noting that the expected number of occurrences of dense extensions is zero.

Definition 6 *Let M and N be L-structures, and let $\mathbf{x} = (x_1, \ldots, x_i)$ and $\mathbf{y} = (y_1, \ldots, y_i)$ be tuples of elements of M and N respectively. We say the tuples are t-equivalent, denoted $\mathbf{x} \equiv_t \mathbf{y}$ if their t-closures, in M and N respectively, are isomorphic with x_i mapped to y_i.*

We conclude this section by showing that equivalent tuples enjoy a nice back-and-forth property. The proof follows the ideas of Theorem 7.3 in [31] but of course needs to be adapted to the new distributions. Note the requirement that $\alpha > 1$, needed in Theorem 4, and that α be irrational, needed in Lemma 1.

Theorem 5 *Let $\alpha > 1$ be irrational, and let $M \sim L_8(n, p)$ and $N \sim P_8(n, p)$ be random 3-CNF formulas drawn from the non-planted and planted distribution respectively, where $p = \Theta(n^{-\alpha})$. For every $b \geq 0$ and $i \geq 0$, there exists $a \geq 0$ such that the following holds almost surely: for every pair of tuples $\mathbf{x} = (x_1, \ldots, x_i)$ in M and $\mathbf{y} = (y_1, \ldots, y_i)$ in N such that $\mathbf{x} \equiv_a \mathbf{y}$, and for every further x_{i+1} in M, there exists y_{i+1} in N such that $(x_1, \ldots, x_{i+1}) \equiv_b (y_1, \ldots, y_{i+1})$, and for every further y_{i+1} in N, there exists x_{i+1} in M such that $(x_1, \ldots, x_{i+1}) \equiv_b (y_1, \ldots, y_{i+1})$.*

Proof: Let K be the constant in Lemma 2 for tuple-length $i + 1$ and closure-bound b. Let $a = b + K$.

Suppose $\mathbf{x} = (x_1, \ldots, x_i)$ and $\mathbf{y} = (y_1, \ldots, y_i)$ are such that $\mathbf{x} \equiv_a \mathbf{y}$. Let $f : cl_a(\mathbf{x}) \to cl_a(\mathbf{y})$ be the isomorphism witnessing that $\mathbf{x} \equiv_a \mathbf{y}$. We consider the forth property, the back is dual. Let $x_{i+1} \in M$ be different from $x_1, \ldots, x_i$. We consider two cases:

Case 1: $x_{i+1} \in cl_K(\mathbf{x})$. Then $cl_b(\mathbf{x}, x_{i+1}) \subseteq cl_a(\mathbf{x})$ by Lemma 1 (i). Set $y_{i+1} = f(x_{i+1})$. Clearly, $cl_b(\mathbf{x}, x_{i+1}) \cong cl_b(\mathbf{y}, y_{i+1})$, so $(x_1, \ldots, x_{i+1}) \equiv_b (y_1, \ldots, y_{i+1})$.

Case 2: $x_{i+1} \notin cl_K(\mathbf{x})$. Let $H = cl_b(\mathbf{x}, x_{i+1})$, and let $S \subseteq H$ be the K-closure of $\mathbf{x}$ in H. Note that $|H| \leq K$ by our choice of K, and that $x_{i+1} \in H - S$. Then, (S, H) is safe by (ii) in Lemma 1. Let (v, e) be its type. Note that $S \subseteq cl_a(\mathbf{x})$. Let $\mathbf{z}$ be an enumeration of S, and let $\mathbf{z}'$ be an enumeration of $f(S)$. By Theorem 4, the number of (S, H)-extensions of $\mathbf{z}'$ is $\Theta(n^{v-\alpha e})$. Therefore, the number of injective homomorphisms $g : cl_b(\mathbf{x}, x_{i+1}) \to M$ with $g(\mathbf{z}) = \mathbf{z}'$ is $\Omega(n^{v-\alpha e})$. By letting $y_{i+1} = g(x_{i+1})$, the images of these homomorphisms are substructures of $cl_b(\mathbf{y}, y_{i+1})$ possibly with missing elements and/or tuples.

We just showed that there are $\Omega(n^{v-\alpha e})$ many y_{i+1} with $cl_b(\mathbf{y}, y_{i+1})$ which is isomorphic to either H or H' for some

H' containing H as a proper substructure. Fix such an H' of the second sort, and let (v', e') be the type of the (S, H') extension that H' forms over S. Note that (H, H') has type $(v' - v, e' - e)$, and if $v < v'$, it is dense because H' is a closure and thus rigid over its roots. Thus, if $v < v'$, then $v' - v - \alpha(e' - e) < 0$. On the other hand, if $v = v'$, then $e < e'$, so $v' - \alpha e' < v - \alpha e$. In both cases we have $v' - \alpha e' < v - \alpha e$. There are at most $O(n^{v'-\alpha e'})$ many (S, H')-extensions of $\mathbf{z}'$, and since the exponent is smaller, this is $o(n^{v-\alpha e})$. This shows that for every fixed possible b-closure containing H as a proper substructure, there are only $o(n^{v-\alpha e})$ y_{i+1} with $cl_b(\mathbf{y}, y_{i+1})$ isomorphic to it. However, b-closures are bounded in size, so there are only a bounded number of possible b-closures up to isomorphism. It follows that the required y_{i+1} exists, and in fact $\Omega(n^{v-\alpha e})$ many. □

The back-and-forth property implied by Theorem 5 can now be used to show that two random 3-CNF formulas drawn from $L_8(n, p)$ and $P_8(n, p)$ are almost surely indistiguishable by first-order formulas of any fixed quantifier rank. The quantifier rank of a formula is defined inductively as follows: $\mathrm{rk}(\varphi) = 0$ if φ is atomic, $\mathrm{rk}(\varphi) = \mathrm{rk}(\psi)$ if $\varphi = \neg\psi$, $\mathrm{rk}(\varphi) = \max\{\mathrm{rk}(\psi), \mathrm{rk}(\theta)\}$ if $\varphi = \psi \wedge \theta$, and $\mathrm{rk}(\varphi) = \mathrm{rk}(\psi) + 1$ if $\varphi = (\exists x)(\psi)$. In other words, $\mathrm{rk}(\varphi)$ is the maximum nesting of quantifiers in φ.

To show this we will use an Ehrenfeucht-Fraïssé game. The game is played by two players, the Spoiler and the Duplicator, on two structures M and N over the same language. In round i of the game, the Spoiler chooses one structure and an element $x_i \in M$ (or $y_i \in N$) of that structure, and the Duplicator replies by choosing one element $y_i \in N$ (or $x_i \in M$) of the other structure. If after r rounds of play, the mapping $x_i \mapsto y_i$ is a partial isomorphism between M and N, then the Duplicator wins the r-round game. Otherwise, the Spoiler wins. The main result about Ehrenfeucht-Fraïssé games is the following:

Theorem 6 ([14]) *Let M and N be L-structures and $r \geq 0$. Then the Duplicator wins the r-round Ehrenfeucht-Fraïssé game on M and N if and only if M and N are indistinguishable by first-order sentences of quantifier rank at most r.*

Now we can show that two random 3-CNF formulas from $L_8(n, p)$ and $P_8(n, p)$ respectively are almost surely a win for the Duplicator. Note again the need for $\alpha > 1$.

Lemma 3 *Let $\alpha > 1$ be irrational, and let $M \sim L_8(n, p)$ and $N \sim P_8(n, p)$ be random 3-CNF formulas drawn from the non-planted and planted distribution respectively, where $p = \Theta(n^{-\alpha})$. For every $r \geq 0$, the Duplicator almost surely wins the r-round Ehrenfeucht-Fraïssé game on M and N.*

Proof: Let $t_r = 0$, and for $i \in \{0, \ldots, r-1\}$ in decreasing order, let t_i be the a in Theorem 5 when $b = t_{i+1}$. Theorem 5 gives then the winning strategy for the Duplicator: at round i, the Duplicator chooses $y_i \in N$ (or $x_i \in M$) in response to the Spoiler's move $x_i \in M$ (or $y_i \in M$) in such a way that $(x_1, \ldots, x_i) \equiv_{t_i} (y_1, \ldots, y_i)$. By the end of the game, we have $(x_1, \ldots, x_r) \equiv_0 (y_1, \ldots, y_r)$, so the Duplicator wins. □

Finally, the proof of Theorem 1:

Proof of Theorem 1.2: Let $p = \Theta(n^{-\alpha})$ for an irrational $\alpha > 1$. Suppose for contradiction that φ implies unsatisfiability and is almost surely true. Let r be the quantifier rank of φ. By Lemma 3 and Theorem 6, two random 3-CNF formulas $M \sim L_8(n, p)$ and $N \sim P_8(n, p)$ are almost surely indistinguishable by φ. Hence, φ holds almost surely on N. However, φ implies unsatisfiability and N is satisfiable. A contradiction. □

Proof of Theorem 1.1: This is the easy part of the theorem. Let $p = \Theta(n^{-\alpha})$ for an arbitrary $0 < \alpha < 1$. Let $K > 0$ be such that $K/(K+1) > \alpha$. Consider the following formula H with K variables:

$$x_1 \wedge (\neg x_1 \vee x_2) \wedge (\neg x_2 \vee x_3) \wedge \cdots \wedge (\neg x_{K-1} \vee x_K) \wedge \neg x_K.$$

Clearly, H is unsatisfiable. By padding each clause with repeated literals we can view it as a 3-CNF formula. The maximal density of its subformulas is $(K+1)/K$. Now we apply Corollary 1 and conclude that M contains a copy of H almost surely. Thus, the subformula induced by the variables of this copy is unsatisfiable. □

6. Conclusion

Our main result establishes the breakpoint where first-order definability can certify unsatisfiability for power probabilities $p = n^{-\alpha}$. We believe the result stands by itself independently of its implications or non-implications about the random 3-SAT hypothesis for polynomial-time. On the one hand, definability on random structures is sometimes surprisingly strong as the motivating example in the introduction shows. On the other hand, definable properties of random 3-CNF formulas, and their correlation with satisfiability or other important properies of formulas, seem to deserve independent study.

Such an independent study would ask about classical concepts of logic on random structures, such as the 0-1 laws, convergence laws, and their deep model-theoretic consequences [5]. Let us mention that it follows from our results that the 0-1 law for sparse random graphs in [29, 30] extends to sparse random 3-CNF formulas, and even to the

planted distribution (always for irrational α). We omit details in this version. One consequence is that the almost sure theory of such random 3-CNF formulas is a complete theory. Another consequence is that our main inexpressibility result is actually stronger. Indeed, it shows that first-order logic cannot even certify unsatisfiability with positive asymptotic probability at $m = \Theta(n^{2-\alpha})$, when $\alpha > 0$ is irrational. Let us note that, somewhat trivially, first-order logic can certify unsatisfiability with inverse polynomial decaying probability. All these questions deserve further study and should be considered somewhere else.

References

[1] R. Aharoni and N. Linial. Minimal unsatisfiable formulas and minimal non two-colorable hypergraphs. *Journal of Combinatorial Theory, Series A*, 43:196–204, 1986.

[2] M. Alekhnovich and E. Ben-Sasson. Linear upper bounds for random walk on small density random 3-CNFs. In *44th Annual IEEE Symposium on Foundations of Computer Science*, pages 352–361, 2003.

[3] A. Atserias. On sufficient conditions for unsatisfiability of random formulas. *Journal of the ACM*, 51(2):281–311, 2004. Preliminary version in LICS 2002.

[4] A. Atserias, M. L. Bonet, and J. L. Esteban. Lower bounds for the weak pigeonhole principle and random formulas beyond resolution. *Information and Computation*, 176(2):136–152, 2002.

[5] J. T. Baldwin and S. Shelah. Randomness and semigenericity. *Transactions of the American Mathematical Society*, 349(4):1359–1376, 1997.

[6] P. Beame, R. Karp, T. Pitassi, and M. Saks. On the complexity of unsatisfiability proofs for random k-CNF formulas. In *30th Annual ACM Symposium on the Theory of Computing*, pages 561–571, 1998.

[7] E. Ben-Sasson and R. Impagliazzo. Random CNF's are hard for the polynomial calculus. In *40th Annual IEEE Symposium on Foundations of Computer Science*, pages 415–421, 1999.

[8] E. Ben-Sasson and A. Wigderson. Short proofs are narrow–resolution made simple. *Journal of the ACM*, 48(2):149–169, 2001.

[9] A. Blass and F. Harary. Properties of almost all graphs and complexes. *Journal of Graph Theory*, 3:225–240, 1979.

[10] J. Buresh-Oppenheim, N. Galesi, S. Hoory, A. Magen, and T. Pitassi. Rank bounds and integrality gaps for cutting planes procedures. In *44th Annual IEEE Symposium on Foundations of Computer Science*, 2003.

[11] V. Chvátal and E. Szemerédi. Many hard examples for resolution. *Journal of the ACM*, 35(4):759–768, 1988.

[12] O. Dubois, Y. Boufkhad, and J. Mandler. Typical random 3-SAT formulae and the satisfiability threshold. In *11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 126–127, 2000.

[13] O. Dubois, R. Monasson, B. Selman, and R. Zecchina (Guest Eds.). Phase transitions in combinatorial problems. *Theoretical Computer Science*, 265(1–2), 2001.

[14] H. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer-Verlag, 1995.

[15] H. Ebbinghaus, J. Flum, and W. Thomas. *Mathematical Logic*. Springer-Verlag, 1984.

[16] P. Erdös and A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci.*, 5:17–61, 1960.

[17] R. Fagin. Probabilities on finite models. *Journal of Symbolic Logic*, 41:50–58, 1976.

[18] U. Feige. Relations between average case complexity and approximation complexity. In *34th Annual ACM Symposium on the Theory of Computing*, pages 534–543, 2002.

[19] J. Friedman and A. Goerdt. Recognizing more unsatisfiable random 3-SAT instances efficiently. In *28th International Colloquium on Automata, Languages and Programming*, volume 2076 of *Lecture Notes in Computer Science*, pages 310–321. Springer-Verlag, 2001.

[20] Y. V. Glebskii, D. I. Kogan, M. I. Liagonkii, and V. A. Talanov. Range and degree realizability of formulas in the restricted predicate calculua. *Cybernetics*, 5:142–154, 1969.

[21] A. Goerdt and M. Krivelevich. Efficient recognition of random unsatisfiable k-SAT instances by spectral methods. In *18th International Symposium on Theoretical Aspects of Computer Science*, volume 2010 of *Lecture Notes in Computer Science*, pages 294–304. Springer-Verlag, 2001.

[22] A. C. Kaporis, L. M. Kirousis, and E. G. Lalas. The probabilistic analysis of a greedy satisfiability algorithm. In *10th Annual European Symposium on Algorithms*, volume 2461 of *Lecture Notes in Computer Science*, pages 574–585. Springer, 2002.

[23] J. H. Kim and V. H. Vu. Concentration of multivariate polynomials and its applications. *Combinatorica*, 20(3):417–434, 2000.

[24] Ph. G. Kolaitis and M. Y. Vardi. 0-1 laws and decision problems for fragments of second-order logic. *Information and Computation*, 87:302–338, 1990.

[25] Ph. G. Kolaitis and M. Y. Vardi. Conjunctive-query containment and constraint satisfaction. *Journal of Computer and System Sciences*, 61(2):302–332, 2000.

[26] J. F. Lynch. Almost sure theories. *Annals of Mathematical Logic*, 18:91–135, 1980.

[27] R. Monasson and R. Zecchina. Statistical mechanics of the random k-satisfiability model. *Physical Review E*, 56(2):1357–1370, 1997.

[28] B. Selman, D. G. Mitchell, and H. J. Levesque. Generating hard satisfiability problems. *Artificial Intelligence*, 81:17–29, 1996.

[29] S. Shelah and J. Spencer. Threshold spectra for random graphs. In *19th Annual ACM Symposium on the Theory of Computing*, pages 421–424, 1987.

[30] S. Shelah and J. Spencer. Zero-one laws for sparse random graphs. *Journal of the American Mathematical Society*, 1(1):97–115, 1988.

[31] J. Spencer. Threshold spectra via the Ehrenfeucht game. *Discrete Applied Mathematics*, 30:235–252, 1991.

[32] J. Spencer. *The Strange Logic of Random Graphs*. Springer, 2001.