Electronic Colloquium on Computational Complexity, Report No. 78 (2026)

Superpolynomial Length Lower Bounds for Tree-Like Semantic Proof Systems with Bounded Line Size

Susanna F. de Rezende
Lund University

David Engström
Lund University

Yassine Ghannane
University of Copenhagen

Kilian Risse
Lund University

May 2, 2026

Abstract

We prove superpolynomial length lower bounds for the semantic tree-like Frege refutation system with bounded line size. Concretely, for any function $n^{2-\varepsilon} \leq s(n) \leq 2^{n^{1-\varepsilon}}$ we exhibit an explicit family $\mathcal{A}$ of n-variate CNF formulas A, each of size $|A| \leq s(n)^{1+\varepsilon}$, such that if A is chosen uniformly from $\mathcal{A}$, then asymptotically almost surely any tree-like Frege refutation of A in line-size $s(n)$ is of length super-polynomial in $|A|$. Our lower bounds apply also to tree-like degree-d threshold systems, for $d \approx \log\big(s(n)\big)$, that is, for d up to $n^{1-\varepsilon}$. More generally, our lower bounds apply to the semantic version of these systems and to any semantic tree-like proof system where the number of distinct lines is bounded by $\exp\big(s(n)\big)$.

ISSN 1433-8092

1 Introduction

Proof complexity studies certificates of unsatisfiability, *refutations*, of unsatisfiable propositional formulas. The original motivation for the conception of the field [Rec75, CR79] was to establish that there are propositional formulas that require refutations of size superpolynomial in the formula size or, equivalently, to separate NP from coNP. Since such a separation would imply P ≠ NP we do not expect a resolution in the near future.

The general approach towards such a separation is to study constrained refutations and show length lower bounds on ever stronger classes of refutations, *proof systems*, to ultimately establish that there is no *polynomially bounded* proof system: the goal is to establish that there is no proof system that has polynomially-size refutations for *all* unsatisfiable propositional formulas. Establishing that there is no polynomially bounded proof system readily separates NP from coNP.

As an intermediary goal proof complexity studies different proof systems and compares their relative deductive abilities. Lower bounds on limited proof systems are interesting in their own right yielding lower bounds for limited models of computation and applications to coding theory [KM24a, KM24b], the theory of total search problems [GHJ+24], and many other adjacent areas. These lower bounds are well-motivated but do not seem to directly contribute towards the original objective of establishing NP ≠ coNP.

Progress on extending lower bounds from weak proof systems such as *resolution*, *cutting planes*, or *bounded-depth Frege* to more powerful systems is slow. If the ultimate aim is to establish lower bounds for strong proof systems that can efficiently refute simple benchmark formulas such as the *pigeonhole principle*, *Tseitin contradictions*, or the *clique-coloring formulas*, then we need to develop lower bound strategies that do not depend on these formulas requiring long refutations in a proof system. Since worst-case refutation size lower bounds on constraint satisfaction problems (CSPs) follow from the hardness of simple combinatorial principles [AO18], it is common to study average-case CSP instances, such as the *clique* or *coloring* formulas on an Erdős-Rényi random graph, and prove refutation size lower bounds on these formulas in weak proof systems. The hope is that since we cannot obtain average-case lower bounds from the hardness of simple formulas, these lower bounds hold regardless of whether simple formulas are hard for a given proof system. While the average-case lower bound arguments for weak proof systems are getting more and more involved, we cannot rule out that these lower bounds hinge on the fact that it is hard to refute simple combinatorial principles.

Our main result is an average-case refutation length lower bound for semantic tree-like proof systems with a bounded number of distinct lines. More concretely we show that there exists an explicit family $\mathcal{A}$ of n-variate CNF formulas A such that if A is chosen uniformly from $\mathcal{A}$, then any tree-like semantic proof system over a bounded number of lines asymptotically almost surely requires super-polynomial length to refute A. These proof systems are very strong, and can efficiently refute simple combinatorial principles such as the pigeonhole principle, the Tseitin contradictions, and the clique-colouring formula, thus establishing that our lower bound technique does not hinge on these principles being hard. As corollaries, we obtain average-case super-polynomial length lower bounds for tree-like Frege refutation systems with lines of bounded size, and for tree-like degree-d threshold proof systems, for d up to $n^{1-\varepsilon}$. Our lower bounds apply for the semantic version of these systems and the parameters are close to optimal: allowing slightly larger line-size, or slightly larger degree, would allow these proof systems to represent the CNF formula A on a single line, thus allowing these semantic systems to immediately derive contradiction.

1.1 Semantic Proof Systems

For concreteness, let us informally define these semantic proof systems; we refer to Section 3 for a formal treatment. Consider a set $\mathcal{F}$ of n-variate Boolean functions. In the following, we view each function $f \in \mathcal{F}$ as an indicator of a set of assignments $\alpha \in \{0,1\}^n$ that are ruled out, i.e., we think of α being ruled out by f if $f(\alpha) = 1$. The *semantic $\mathcal{F}$ proof system*, denoted by $\mathrm{sem}(\mathcal{F})$, is an inferential proof system that operates over proof lines in $\mathcal{F}$: a $\mathrm{sem}(\mathcal{F})$ *refutation* π of a CNF formula $A \coloneqq C_1 \wedge \cdots \wedge C_m$ is a sequence $\pi \coloneqq (f_1, \ldots, f_t)$ such that $f_t = 1$ is the constant 1 function and each f_i is either an axiom, that is, there is a clause $C \in A$ such that f is the indicator of falsifying assignments of C; or $f_i \in \mathcal{F}$ and there are f_j, f_k with $j, k < i$ such that f_i is (semantically) implied by f_j and f_k, that is, for all $\alpha \in \{0,1\}^n$ if $f_i(\alpha) = 1$ then either $f_j(\alpha) = 1$ or $f_k(\alpha) = 1$. The *length* of the derivation π is t, and π is *tree-like* if each function $f_j \in \pi$ is used at most once to derive some other function f_i. We denote the tree-like $\mathrm{sem}(\mathcal{F})$ proof system by $\mathrm{sem}^\star(\mathcal{F})$.

Semantic tree-like proof systems are very powerful: for any fixed CNF formula $A \coloneqq C_1 \wedge \cdots \wedge C_m$ there is a set $\mathcal{F}$ of Boolean functions such that there are semantic refutations of A over $\mathcal{F}$ of size $O(m)$: if $\mathcal{F}$ contains the functions $\left\{f_i \coloneqq \bigwedge_{j=1}^{i} C_j \,\middle|\, i \in [m]\right\}$, then there is a $\mathrm{sem}^\star(\mathcal{F})$ refutation of A by iteratively deriving the functions f_i for $i = 1, \ldots, m$. More generally, even some large formula families $\mathcal{A}$ may be refuted by $\mathrm{sem}^\star(\mathcal{F})$ for small sets $|\mathcal{F}| \ll |\mathcal{A}|$. For example, the family of Tseitin contradictions on at most n variables may be refuted by adding all linear equations mod 2 to $\mathcal{F}$, and the graph pigeonhole formulas over at most n variables may be refuted by adding linear inequalities to $\mathcal{F}$.

Furthermore, if $|\mathcal{F}| = \exp\big(\Omega(n \log n)\big)$ we can assume—without increasing the size of $\mathcal{F}$ by more than a polynomial factor—that $\mathrm{sem}^\star(\mathcal{F})$ is well-behaved in the sense that it is complete and closed under permutations of variables. More formally, for any $\mathcal{F}$, we can construct a $\mathcal{F}' \supseteq \mathcal{F}$ of size at most $|\mathcal{F}| \cdot \exp\big(O(n \log n)\big)$ that is well-behaved. Indeed, it suffices to include the 3^n functions representing all clauses, to ensure that $\mathrm{sem}^\star(\mathcal{F}')$ is complete. In order to make it closed under permutations, for every $f \in \mathcal{F}$, we further include f under all permutations of variables. This increases the size by at most a multiplicative factor $\exp\big(O(n \log n)\big)$.

We can therefore conclude that there are sets $\mathcal{F}$ of size $\exp\big(O(n \log n)\big)$ such that $\mathrm{sem}^\star(\mathcal{F})$ is a well-behaved proof system, and it refutes any fixed formula and even simple families of formulas. In general, these proof systems can be somewhat unnatural. For instance, we cannot guarantee that they are closed under restrictions nor that proofs are efficiently verifiable. In this paper, we consider concrete, natural proof systems that are captured by this definition of semantic proof systems. We highlight two of these, for which, prior to this paper, no superpolynomial refutation size lower bounds were known.

The first of these are tree-like Frege refutation systems operating over Boolean formulas of size at most s. We consider the refutation version, where the system is given a CNF formula, viewed as a set of axioms, and every line of the refutation is either one of these axioms or can be derived by some sound inference rule from two previous lines. This system can be simulated by $\mathrm{sem}^\star(\mathcal{F}_s)$, where $\mathcal{F}_s$ is the family of formulas of size at most s, and thus $|\mathcal{F}_s| \leq \exp\big(O(s \log s)\big)$. We note that while tree-like Frege simulates general, dag-like Frege, it is unclear whether this also holds for Frege systems with bounded line-size. This is because, given a length-ℓ dag-like Frege refutation of a CNF formula A with lines containing formulas of size at most s, the known simulation produces a tree-like Frege proof of length $O(\ell)$ but increases the maximum size of the formulas in a line to $\ell\, s$. This might seem like a mild increase in formula size but, in the semantic setting, allowing for formulas of size $|A|$ yields trivial upper bounds—the system can

represent the formula A in one line and thus immediately derive contradiction.

The second proof system we highlight are tree-like degree-d threshold systems which operate over polynomial inequalities of degree at most d. Since there are $\exp\big(O(n^{d+1})\big)$ many distinct such inequalities, there is a family $\mathcal{F}$ of functions, with $|\mathcal{F}| \leq \exp\big(O(n^{d+1})\big)$ such that $\mathrm{sem}^\star(\mathcal{F})$ simulates tree-like degree-d threshold systems. We note that there are known superpolynomial lower bounds for this system when d is at most logarithmic in n [BPS07, GP18, IR21]. Similarly to the case of Frege, once the semantic version of these systems can encode a contradictory CNF formula A in one line, it can immediately derive contradiction. This implies in particular that semantic tree-like degree-d threshold systems can refute any d-CNF formula in length linear in the formula size.

1.2 Lower Bounds for Tree-Like Semantic Proof Systems

In order to prove refutation length lower bounds on $\mathrm{sem}^\star(\mathcal{F})$ refutations that hold for *any* set $\mathcal{F}$ of bounded size, we need to consider a family of formulas $\mathcal{A}$ of size $|\mathcal{A}| \gg |\mathcal{F}|$ and argue that if A is sampled uniformly from $\mathcal{A}$, then asymptotically almost surely there are no short tree-like $\mathrm{sem}(\mathcal{F})$ refutations of A. We present our lower bounds in three different parameter settings. The first lower bound minimizes the width of the family of CNF formulas $\mathcal{A}$.

Theorem 1.1 (informal)**.** *There is a family $\mathcal{A}$ of n-variate 3-CNF formulas such that for* any *set $\mathcal{F}$ of n-variate Boolean functions of size $|\mathcal{F}| \leq \exp(n^{2-\varepsilon})$ the following holds. If A is sampled uniformly from $\mathcal{A}$, then asymptotically almost surely all* $\mathrm{sem}^\star(\mathcal{F})$ *refutations of A are of super-polynomial length.*

As corollaries we obtain optimal average-case lower bounds for the tree-like versions of cutting planes and resolution over linear equations proof systems. We discuss these consequences in Section 1.3.

Note that the 3-CNF formulas $A \in \mathcal{A}$ for which Theorem 1.1 holds have to be rather dense since the set $\mathcal{A}$ needs to be of size $|\mathcal{A}| \gg |\mathcal{F}|$, for any $\mathcal{F}$ of size $\exp\big(O(n^{2-\varepsilon})\big)$. If we want to allow the proof system to operate over significantly more lines, say, $|\mathcal{F}| \gg \exp(n^3)$, then our family $\mathcal{A}$ of hard formulas needs to be over formulas of larger width as there are only $O\big(\exp(n^3)\big)$ CNF formulas of width 3. The following statement is the intermediate parameter setting for CNF formulas of constant width.

Theorem 1.2 (informal)**.** *For any constant even integer $\ell \geq 4$, there is a family $\mathcal{A}$ of n-variate ℓ-CNF formulas such that for* any *set $\mathcal{F}$ of n-variate Boolean functions of size $|\mathcal{F}| \leq \exp\big(n^{\ell-\varepsilon}\big)$ the following holds. If A is sampled uniformly from $\mathcal{A}$, then asymptotically almost surely all tree-like* $\mathrm{sem}(\mathcal{F})$ *refutations of A are of super-polynomial length.*

Note that the size upper bound on $\mathcal{F}$ is essentially optimal as there are $\exp\big(O(n^\ell)\big)$ many ℓ-CNF formulas. Recall that for the theorem to hold it must be the case that $|\mathcal{A}| \gg |\mathcal{F}|$ and hence most formulas in $\mathcal{A}$ must be dense. For the family $\mathcal{A}$ we consider, most formulas $A \in \mathcal{A}$ are of size $\widetilde{\Theta}(n^\ell)$.

Theorem 1.2, moreover, establishes a strict hierarchy of tree-like semantic proof systems

$$\mathrm{sem}^\star(\mathcal{F}_4) \subsetneq \mathrm{sem}^\star(\mathcal{F}_6) \subsetneq \mathrm{sem}^\star(\mathcal{F}_8) \subsetneq \cdots \tag{1}$$

where $\mathrm{sem}^\star(\mathcal{F}_\ell)$ refutes *any* k-CNF formula F with $k < \ell$ in length linear in $|F|$ but, at the same time, cannot refute all ℓ-CNF formulas F in length polynomial in $|F|$. This implies that this hierarchy is strict with respect to refutations of polynomial length.

Finally, we observe that both theorems above imply that most minimially unsatifiable constant width CNF formulas are dense. In particular, Theorem 1.1 shows that most minimally unsatisfiable 3-CNF formulas have at least $n^{2-\varepsilon}$ clauses, and Theorem 1.2 shows that for even $\ell \geq 4$ most minimally unsatisfiable ℓ-CNF formulas are of almost maximum density. This finding was somewhat surprising to us, although it is possible that it has already been established previously by other methods.

The final parameter regime we consider optimizes the number of lines $\mathcal{F}$ while maintaining that the lower bound is super-polynomial in the formula size.

Theorem 1.3 (informal)**.** *For every function $s(n)$, satisfying $n^{2-\varepsilon} \leq s(n) \leq 2^{n^{1-\varepsilon}}$ we exhibit an explicit family $\mathcal{A}$ of n-variate CNF formulas A, each of size $|A| \leq s(n)^{1+\varepsilon}$, such that for* any *set $\mathcal{F}$ of n-variate Boolean functions of size $|\mathcal{F}| \leq \exp\big(s(n)\big)$ the following holds. If A is sampled uniformly from $\mathcal{A}$, then asymptotically almost surely all tree-like* $\mathrm{sem}(\mathcal{F})$ *refutations of A are of length superpolynomial in $|A|$.*

As before, in order to prove lower bounds for tree-like semantic $\mathcal{F}$ proof systems for larger families $\mathcal{F}$, we need to consider CNF formulas A of larger width and over more clauses. In this sense, we can allow for larger families $\mathcal{F}$ in terms of the number of variables of A, but not in terms of its size. Nevertheless, the lower bound we obtain is a length lower bound, that is, a lower bound in the number of steps in the refutation, and it is superpolynomial in the number of clauses of A. We note, moreover, that it is even possible to push the size of $\mathcal{F}$ to $\exp\big(\exp\big(O(n)\big)\big)$ albeit at the expense of the lower bound: it becomes polynomial in the formula size instead of super-polynomial.

We establish Theorems 1.1 to 1.3 by extending the lower bound approach of constructing a *pseudo-measure* [dRPR23]. This approach has been successfully employed to show essentially optimal average-case clique refutation size lower bounds for the Sherali–Adams proof system with bounded coefficients [dRPR23]. It has also proven useful to analyse TFNP intersection classes [HKT24] and similar ideas were used to obtain total coefficient size lower bounds on Nullstellensatz refutations [PZ24].

The very high-level proof idea of [dRPR23] adapted to our setting is to construct a linear pseudo-measure μ that assigns contradiction to 1 but any weakening of an axiom to a value of small magnitude. Since the measure is linear, the measure of all the leaves of a tree-like proof, that is, weakenings of axioms, need to sum to the measure of contradiction, which is equal to 1. Since the measure assigns each leaf small value, there must be many leaves.

In the setting of [dRPR23], all the weakenings of axioms are very structured. This is certainly not the case in our setting as we allow arbitrary proof lines without any syntactic restriction. One of our contributions is to extend their proof strategy to any fixed set of weakenings of axioms regardless of their structure. At the heart of the argument is a union bound over all possible weakenings of axioms of any refutation. Since we have a bound on the number of distinct proof lines we also have a bound on the number of such weakenings of axioms which allows us to appeal to the union bound. This argument is in some sense a delicate counting argument. Finally, let us mention that we do *not* rely on the non-negativity argument of [dRPR23], which does seem to rely on the precise structure of proof lines.

As mentioned in Section 1.1, the set $\mathcal{F}$ of proof lines may be chosen such that $\mathrm{sem}^\star(\mathcal{F})$ refutes *all* standard benchmark formulas based on simple combinatorial principles. Hence the approach of constructing a pseudo-measure may have the potential to yield lower bounds for even stronger proof systems.

Corollary 1.4. *Lower bounds obtained by the pseudo-measure approach do not rely on the fact that simple combinatorial principles are hard for a given proof system.*

While Theorem 1.1 highlights a major strength of the pseudo-measure approach it at the same time points to its main weakness: in its current form the approach does not depend on the precise syntactic rules of a proof system. Whether this approach can be adapted so that it depends on the precise syntactic derivation rules is an interesting direction, which we leave as an open problem.

The family of formulas for which we exhibit the superpolynomial lower bound of Theorems 1.1 to 1.3 is the family of *k-clique formulas*. For each graph $G \in \{0,1\}^{\binom{n}{2}}$ this family contains a formula claiming that G contains a k-clique. It is well-known that the maximum clique size of a graph sampled by including each edge independently with probability $1/2$ is bounded by $2(1+o(1))\log n$. Hence formulas sampled uniformly from this family are with high probability unsatisfiable for $k \geq 3\log n$. Note that since this family is of size $2^{\binom{n}{2}} = \exp(O(n^2))$ the restriction on the size of F in Theorem 1.1 is essentially tight. To obtain Theorem 1.2 we consider non-standard encodings of the clique formula that interpolate between the standard unary encoding used for Theorem 1.1 and the harder to refute binary encoding used to obtain Theorem 1.3. To get the full range of parameters in Theorem 1.3, apart from the different encodings, we also need to consider smaller edge probability so that there are no k-cliques of size $\log^{\varepsilon} n$.

1.3 Lower Bounds for Tree-Like Cook–Reckhow Proof Systems

From Theorem 1.3, we obtain the first superpolynomial refutation length lower bounds on tree-like $\mathrm{Th}(d)$ refutations, i.e., refutations with lines consisting of degree-d polynomial inequalities, for d polynomial in the number of variables. The previously strongest lower bounds held for d logarithmic in the number of variables [GP18, IR21], both building on connections to d-party communication complexity [BPS07].

Corollary 1.5 (informal)**.** *There is a family $\mathcal{A}$ of n-variate formulas such that if A is sampled uniformly from $\mathcal{A}$, then asymptotically almost surely any tree-like* $\mathrm{Th}\left(n^{1-\varepsilon}\right)$ *refutation of A is of length superpolynomial in $|A|$.*

Similarly, we obtain length lower bounds for tree-like Frege refutations with lines of bounded size.

Corollary 1.6 (informal)**.** *For every function $n^{2-\varepsilon} \leq s(n) \leq 2^{n^{1-\varepsilon}}$, there is a family $\mathcal{A}$ of n-variate CNF formulas A, each of size $|A| \leq s(n)^{1+\varepsilon}$, such that the following holds. If A is sampled uniformly from $\mathcal{A}$, then asymptotically almost surely any tree-like Frege refutation of A in line-size $s(n)$ is of length super-polynomial in $|A|$.*

As discussed earlier, the result holds for semantic tree-like Frege with line-size at most $s(n)$ and, in terms of allowed line-size, our parameters are close to optimal: allowing line-size $s(n)^{1+\varepsilon} \geq |A|$ would give trivial refutations of A in semantic tree-like Frege.

Corollaries 1.5 and 1.6 hold for k-clique formulas in the binary encoding. Regarding the unary encoding, as consequences of Theorem 1.1, we obtain essentially optimal average-case clique size lower bounds for tree-like versions of the well-studied proof systems cutting planes and resolution over parities.

Corollary 1.7 (informal)**.** *If $G \sim \mathcal{G}(n, 1/2)$ is an Erdős-Rényi random graph, then it holds asymptotically almost surely that tree-like cutting planes and tree-like resolution over parities require length $n^{\Omega(\log n)}$ to refute that G contains a clique of size $n^{\Omega(1)}$.*

Organization. In Section 2 we recall some preliminaries followed by Section 3 in which we formally introduce semantic proof systems. This is followed by Section 4 in which we state our main theorem and prove the consequences mentioned in the introduction. Section 5 is devoted to the proof of our main theorem. We provide some concluding remarks in Section 6.

2 Preliminaries

All logarithms are base 2, for integer $n \in \mathbb{N}^+$ we introduce the shorthand $[n] = \{1, \ldots, n\}$, and sometimes identify singletons $\{u\}$ with the element u. For a set S denote by $\binom{S}{\ell}$ the family of subsets of S of size ℓ and, for a random variable X and an event P, let $\mathbb{1}_P(X)$ be the indicator random variable that is 1 if P holds and 0 otherwise.

2.1 Graph Theory

For $n \in \mathbb{N}$ and $p \in [0,1]$ let $\mathcal{G}(n,p)$ denote the Erdős-Rényi random graph distribution over n-vertex graphs where each of the $\binom{n}{2}$ edges is included independently with probability p. It is well-known that for $p \leq n^{-2/k}$ graphs $G \sim \mathcal{G}(n,p)$ asymptotically almost surely do not contain a k-clique.

Unless stated otherwise, for the remainder of this paper $G = (V,E)$ always denotes a k-partite graph with partition $V = \bigsqcup_{i=1}^{k} V_i$, where each *block* V_i is of size $|V_i| = n$. A k-partite graph $G \sim \mathcal{G}(n,k,p)$ is sampled by first sampling $G' \sim \mathcal{G}(nk,p)$ and intersecting G' with the complete k-partite graph over blocks of size n.

2.2 Extremal Set Theory

Denote by $\mathcal{F}$ a family of sets, and say that a set X is *shattered* by $\mathcal{F}$ if $X \cap \mathcal{F} = 2^X$, for $X \cap \mathcal{F} \coloneqq \{F \cap X \mid F \in \mathcal{F}\}$. The *VC-dimension* of $\mathcal{F}$, denoted by $\mathrm{VC}(\mathcal{F})$, is the largest integer d such that there is a set $X \subseteq \bigcup_{F \in \mathcal{F}} F$ of size $d = |X|$ shattered by $\mathcal{F}$. We rely on the following theorem relating the VC-dimension of a family to its size.

Theorem 2.1 (Sauer-Shelah [Sau72, She72])**.** *For a family of sets $\mathcal{F}$ with $\mathrm{VC}(\mathcal{F}) = d$ and $n \coloneqq \left|\bigcup_{F \in \mathcal{F}} F\right|$ it holds that $|\mathcal{F}| \leq \sum_{i=0}^{d} \binom{n}{i} \leq n^d$.*

2.3 Proof Complexity

All variables considered are Boolean over the standard $\{0,1\}$ basis. A *literal* ℓ is either a variable x or its negation $\bar{x}$, a *clause* C is a disjunction of literals over distinct variables $C \coloneqq \ell_1 \vee \cdots \vee \ell_k$, and a *CNF formula* $F \coloneqq C_1 \wedge \cdots \wedge C_m$ is a conjunction of clauses. We sometimes refer to the clauses of a formula F by *axioms*, write $\mathrm{Vars}(F)$ for the variables of F, and for a Boolean function $f\colon \{0,1\}^n \to \{0,1\}$ let $\mathrm{CNF}(f)$ denote the canonical CNF encoding of f.

A *semantic inferential refutation system* P is defined over a set of proof lines $\mathcal{L}$, and a *P-derivation* π of $L \in \mathcal{L}$ from a set of axioms $A \subseteq \mathcal{L}$ is a sequence $\pi \coloneqq (L_1, \ldots, L_s)$ such that $L_s = L$ and every line $L_i \in \mathcal{L}$ is either an axiom, that is, $L_i \in A$, or it is a semantic consequence over Boolean assignments $\{0,1\}^n$ of at most two previous lines of π. A *P-refutation* π from a set of axioms $A \subseteq \mathcal{L}$ is a P-derivation of contradiction, the *length* of a derivation is the number of lines in it, and a derivation is *tree-like* if every line is used at most once as the premise of an inference and otherwise called *dag-like*.

The *semantic cutting planes* (CP) proof system is an inferential refutation system that operates over linear inequalities with integer coefficients; contradiction is represented by $0 \geq 1$ and a CNF formula is expressed as a system of linear inequalities by translating each of its clauses $C := \bigvee_{i \in I} x_i \vee \bigvee_{j \in J} \bar{x}_j$ to an inequality $\sum_{i \in I} x_i + \sum_{j \in J} (1 - x_j) \geq 1$. *Semantic degree-d threshold proof systems* [BPS07] are a natural generalization of cutting planes: for $d \in \mathbb{N}$, the semantic threshold proof system $\mathrm{Th}(d)$ is the semantic inferential refutation system whose proof lines are polynomial inequalities $p \geq 0$ for $p \in \mathbb{Z}[x_1, \ldots, x_n]$ of degree at most d.

The *semantic Frege* refutation system is a semantic inferential refutation system operating over Boolean formulas, and *semantic Frege with bounded line size s* operates over Boolean formulas of size at most s. For a family of Boolean functions $\mathcal{F}$ closed under negation, that is, if $f \in \mathcal{F}$, then $\neg f \in \mathcal{F}$, let an $\mathcal{F}$-clause $f_1 \vee \cdots \vee f_w$ denote a disjunction of formulas $f_i \in \mathcal{F}$. *Semantic resolution over* $\mathcal{F}$ is the semantic inferential refutation system over $\mathcal{F}$-clauses. Ordinary resolution is recovered by considering the set of Boolean functions that depend on single variables and *resolution over parities* [IS20], denoted $\mathrm{Res}(\oplus)$, is resolution over $\mathbb{F}_2$ affine linear forms.

2.4 Clique Formulas

Before discussing k-partite graphs, let us consider an ordinary graph $G = (V, E)$ over n vertices. The unary k-clique formula over G, denoted by $\mathrm{Clique}(G, k)$, is defined over variables $\{x_{v,i} \,|\, v \in V \text{ and } i \in [k]\}$ where a variable $x_{v,i}$ indicates whether v is the ith clique member. The axioms of $\mathrm{Clique}(G, k)$ are

$$\bar{x}_{u,i} \vee \bar{x}_{v,j} \qquad \forall\, i \neq j \in [k],\ \forall\, \{u, v\} \notin E \qquad \text{(edge axiom)}$$

$$\bigvee_{v \in V} x_{v,i} \qquad \forall\, i \in [k]\ , \qquad \text{(block axiom)}$$

where the edge axioms ensure that two non-adjacent vertices are not simultaneously clique members and the block axioms ensure that at least one vertex is the ith clique member. It should be evident that the formula is satisfiable if and only if G contains a k-clique.

Block Encodings Consider a k-partite graph $G = (V, E)$ with blocks $V_1, \ldots, V_k$ of size n each. We want to encode the claim that G contains a k-clique. Note that if there is a k-clique in G, then it contains exactly one vertex v_i from each block V_i. In other words, the only vertex sets $t \subseteq V$ that *might* be a k-clique satisfy $t \cap V_i = \{v_i\}$ for all $i \in [k]$. Denote the family of these sets by $\mathcal{T} := \prod_{i \in [k]} V_i$ and call a set $t \in \mathcal{T}$ a *tuple*.

Let Y be a set of Boolean variables of size $|Y| = \binom{k}{2} n^2$ and think of these variables as encoding a k-partite graph with blocks of size n each. For a set X of Boolean variables consider any CNF formula $\mathrm{BlockClique}\colon \{0,1\}^X \times \{0,1\}^Y \to \{0,1\}$ mapping k-partite graphs G and assignments to X such that (1) each clause $C \in \mathrm{BlockClique}$ contains at most one literal over Y, (2) if G does *not* contain a k-clique, then for all assignments $\rho \in \{0,1\}^X$ it holds that $\mathrm{BlockClique}(G, \rho) = 0$, and (3) for each tuple $t \in \mathcal{T}$ there is a "witnessing" assignment $\rho_t \in \{0,1\}^X$ such that for all graphs $G \in \{0,1\}^Y$ it holds that

$$t \text{ is a } k\text{-clique in } G \Leftrightarrow \mathrm{BlockClique}(G, \rho_t) = 1\ . \tag{2}$$

In other words, there is a one-to-one correspondence between witnessing assignments ρ_t and tuples $t \in \mathcal{T}$. Our results hold for any encoding satisfying the above provided X is not too large $|X| \leq n^{2-\varepsilon}$. As parameters crucially depend on the size of X let us discuss three concrete encodings.

Unary Block Encoding For a k-partite graph $G = (V, E)$ over blocks $V_1, \dots, V_k$ of size n each let $\text{BlockClique}^{\text{U}}(G, k)\colon \{0,1\}^X \to \{0,1\}$ be the CNF formula defined over variables $X \coloneqq \{x_v \mid v \in \bigcup_i V_i\}$ and consisting of clauses

$$\bar{x}_u \vee \bar{x}_v \qquad \forall \{u, v\} \notin E \qquad \text{(edge axiom)}$$

$$\bigvee_{v \in V_i} x_v \qquad \forall i \in [k] \ . \qquad \text{(block axiom)}$$

The formula $\text{BlockClique}^{\text{U}}(G, k)$ is satisfiable if and only if there is a tuple $t \in \mathcal{T}$ such that the vertex induced subgraph $G[t]$ is a k-clique. Refuting the unary block encoding is at least as hard as refuting the natural (non-block) encoding as summarized by the next statement.

Proposition 2.2 ([BIS07]). *Let $k, n \in \mathbb{N}^+$ be integer and let G be an ordinary graph on kn vertices. If a proof system is closed under restrictions, then the minimum refutation length to refute the* $\text{Clique}(G, k)$ *formula is bounded from below by the minimum length required to refute* $\text{BlockClique}^{\text{U}}(G, k)$ *with respect to any k-partition of G.*

In light of Proposition 2.2 we only consider the block encoding of the k-clique formula. For Theorems 1.2 and 1.3 we need to consider formulas using fewer variables than the unary clique formula.

Binary Block Encoding The binary encoding of the k-clique formula $\text{BlockClique}^{\text{B}}(G, k)$ is defined over variables $X \coloneqq \left\{x_{i,b} \,\middle|\, i \in [k] \text{ and } b \in \left[\lceil \log n \rceil\right]\right\}$. For $a \in [n]$ with binary expansion $a_1 \dots a_{\lceil \log n \rceil}$ denote by $C_{i,a} \coloneqq \bigvee_{b:a_b=0} x_{i,b} \ \vee \bigvee_{b:a_b=1} \bar{x}_{i,b}$ the clause that is falsified if and only if the variables of block i encode a. The formula $\text{BlockClique}^{\text{B}}(G, k)$ consists of axioms

$$C_{i,a} \vee C_{j,b} \qquad \forall\, i \neq j \in [k],\ \forall a, b \text{ such that } \{v_{i,a}, v_{j,b}\} \notin E(G) \ . \qquad \text{(edge axiom)}$$

It should be clear that $\text{BlockClique}^{\text{B}}(G, k)$ is satisfiable if and only if $\text{BlockClique}^{\text{U}}(G, k)$ is. The main reason to consider $\text{BlockClique}^{\text{B}}(G, k)$ is that for $k = O(\log n)$ the number of variables is exponentially smaller than the size of the formula.

Interpolated Block Encodings We can interpolate between the unary and the binary block formulas as follows. Let $m < n$ be a positive integer and suppose for the sake of exposition that $n = m^c$ for integer c (otherwise take $c = \lceil \log(n)/\log(m) \rceil$ and pad). For each $i \in [k]$, write $V_i = \left\{v_{i,\alpha} \,\middle|\, \alpha = (\alpha_1, \dots, \alpha_c) \in \{0, \dots, m-1\}^c\right\}$. The c-ary blocked k-clique formula, denoted $\text{BlockClique}_c(G, k)$, is defined over variables

$$X \coloneqq \left\{x_{i,j,a} \,\middle|\, i \in [k],\ j \in [c],\ a \in \{0, \dots, m-1\}\right\} \ , \tag{3}$$

where $x_{i,j,a}$ indicates that the jth coordinate of the vertex chosen from V_i is a. For a vertex $v_{i,\alpha} \in V_i$, define the clause $C_{v_{i,\alpha}} \coloneqq \bigvee_{j=1}^{c} \bar{x}_{i,j,\alpha_j}$, so that $C_{v_{i,\alpha}}$ is falsified if and only if the mentioned variables encode vertex $v_{i,\alpha}$. The formula $\text{BlockClique}_c(G, k)$ consists of clauses

$$C_u \vee C_v \qquad \forall \{u, v\} \notin E(G) \qquad \text{(edge axiom)}$$

$$\bigvee_{a=0}^{m-1} x_{i,j,a} \qquad \forall i \in [k],\ \forall j \in [c] \ . \qquad \text{(range axiom)}$$

Note that the formula $\text{BlockClique}_1(G, k)$ corresponds to the unary clique encoding whereas the formula $\text{BlockClique}_{\lceil \log n \rceil}(G, k)$ almost corresponds to the binary encoding: the only (minor) difference is that the interpolated encoding has two variables per coordinate of a vertex (of which precisely one is set to 1), whereas the binary encoding has a single variable per coordinate.

2.5 $\mathcal{F}$-Decision Trees

A *decision tree* is a directed tree where each node has in-degree 1 except the designated *root node* which has in-degree 0. The out-degree of each node is either 2 (an *internal node*) or 0 (a *leaf node*), internal nodes are labelled with a variable, and one out-edge of each internal node is labelled 0 whereas the other is labelled 1. Given a Boolean assignment ρ we evaluate a decision tree T by starting at the root node v, and repeatedly considering the label x of the current node and following the edge labelled $\rho(x)$ until we reach a leaf.

For a fixed unsatisfiable CNF formula $A \coloneqq C_1 \wedge \cdots \wedge C_m$ the *falsified clause search problem* Search(A) asks for a clause $C \in A$ falsified by the provided assignment $\rho \in \{0,1\}^{\mathrm{Vars}(A)}$. A *decision tree for* Search(A) is a decision tree with leaves labelled by clauses of A such that any Boolean assignment ρ ending in a leaf labelled $C \in A$ falsifies the clause C, that is, $C\restriction_\rho = 0$.

More generally, we may consider an *$\mathcal{F}$-decision tree*, for $\mathcal{F}$ a set of Boolean functions, defined as a decision tree with internal nodes labelled by functions $f \in \mathcal{F}$. Given an assignment ρ the evaluation proceeds similarly by repeatedly considering the label f of the current node and following the edge labelled $f(\rho)$ until a leaf is reached. An *$\mathcal{F}$-decision tree solves* Search(A) if every leaf ℓ is labelled by a clause $C \in A$ such that every assignment ρ reaching ℓ falsifies C.

3 Semantic Proof Systems Over a Bounded Set of Lines

This section is devoted to a rigorous treatment of the semantic proof systems mentioned in Section 1.1. The main difference to the informal discussion in the introduction is that we consider proof lines as sequences for a parameter $n \in \mathbb{N}$.

Definition 3.1 (Semantic Proof System)**.** Let $\mathcal{F} \coloneqq (F_n \colon n \in \mathbb{N})$ be a sequence of families F_n of n-variate Boolean functions. The *semantic $\mathcal{F}$ proof system*, denoted by sem$(\mathcal{F})$, is an inferential proof system operating over lines $\mathcal{F}$; a sem$(\mathcal{F})$ *refutation* π of an n-variate CNF formula A is a sequence $\pi \coloneqq (f_1, \ldots, f_t)$ such that $f_t = 1$ is the constant 1 function and each f_i is either

- an axiom, that is, there is a clause $C \in A$ such that for all assignments $\alpha \in \{0,1\}^n$ it holds that C is not satisfied by α if and only if $f_i(\alpha) = 1$, or
- $f_i \in F_n$ and there are f_j, f_k with $j, k < i$ such that $f_i \subseteq f_j \cup f_k$ when viewed as indicators.

The *length* of the derivation π is t, and π is *tree-like* if each function $f_j \in \pi$ is used at most once to derive some other function f_i. We denote the tree-like semantic $\mathcal{F}$ proof system by sem$^\star(\mathcal{F})$.

Not every sequence $\mathcal{F}$ gives rise to a proof system that is well-behaved. In particular, the proof system sem$(\mathcal{F})$ is generally not even complete and is not closed under restrictions. In the following we show that for each $\mathcal{F}$ there is a slightly larger $\mathcal{F}'$ that is complete, closed under permutations of variables, and monotone.

Definition 3.2 (Reasonable Sequence)**.** A sequence $\mathcal{F} \coloneqq (F_n \colon n \in \mathbb{N})$ is *reasonable* if the proof system sem$(\mathcal{F})$ is complete and for every $n \in \mathbb{N}$ and every function $f_n \in F_n$ it holds that

1. the function $f_n^\sigma(x_1, \ldots, x_n) \coloneqq f_n(x_{\sigma(1)}, \ldots, x_{\sigma(n)})$ is in F_n, for any permutation $\sigma : [n] \to [n]$, and (closed under permutation)
2. the function $f_{n+1}(x_1, \ldots, x_n, x_{n+1}) \coloneqq f_n(x_1, \ldots, x_n)$ belongs to F_{n+1}. (monotonicity)

It is not too hard to see that any sequence $\mathcal{F}$ gives rise to a reasonable sequence $\mathcal{F}'$ of similar size as stated next.

Lemma 3.3. *For any sequence $(F_n : n \in \mathbb{N})$ there is a reasonable sequence $(F'_n : n \in \mathbb{N})$ such that for all $n \in \mathbb{N}$ it holds that $F_n \subseteq F'_n$ and that $|F'_n| \leq \left(1 + \sum_{i \leq n} |F_i|\right) \cdot \exp\big(O(n \log n)\big)$.*

Proof. By adding all 3^n clauses over n variables to each F'_n we ensure that $\mathrm{sem}(F'_n : n \in \mathbb{N})$ is complete. By further adding for each $f \in F_i$ and all $i \leq n$ the formula $f'(x_1, \ldots, x_n) \coloneqq f(x_1, \ldots, x_i)$ to F'_n we ensure monotonicity. Finally, we can close the sets of functions under permutations by increasing the size by at most a multiplicative factor $\exp\big(O(n \log n)\big)$. □

In the following we study tree-like $\mathrm{sem}(\mathcal{F})$ proof systems where the sequence $\mathcal{F} \coloneqq (F_n : n \in \mathbb{N})$ is only restricted by the size of the families of Boolean functions F_n. To ensure that these systems can be chosen to be somewhat natural we want the families F_n to be of size at least $|F_n| = \exp\big(\Omega(n \log n)\big)$. We say that a family of formulas is "simple" or "based on a simple combinatorial principle" if the entire family can be refuted by a semantic proof system over a small number of lines.

Definition 3.4 (Simple)**.** A sequence $(\mathcal{A}_n : n \in \mathbb{N})$ of families $\mathcal{A}_n$ of n-variate CNF formulas is *simple* if there is a sequence $\mathcal{F} \coloneqq (F_n : n \in \mathbb{N})$ of families F_n of n-ary Boolean functions of size $|F_n| \leq \exp\big(O(n \log n)\big)$ such that for all $n \in \mathbb{N}$ tree-like $\mathrm{sem}(\mathcal{F})$ refutes any formula $A \in \mathcal{A}_n$ in size polynomial in $|A|$.

Before we can discuss concrete examples of simple families we need to discuss when we can translate between different encodings of these formulas.

Definition 3.5 (Encoding Robust)**.** A sequence $(G_n : n \in \mathbb{N})$ of families G_n of n-ary Boolean functions is *encoding robust* if there is a sequence $\mathcal{F} \coloneqq (F_n : n \in \mathbb{N})$ of families F_n of n-ary Boolean functions of size $|F_n| \leq \exp\big(O(n \log n)\big)$ such that for any $n \in \mathbb{N}$ and any $g \in G_n$ it holds that there is a tree-like $\mathrm{sem}(\mathcal{F})$ derivation of g from $\mathrm{CNF}(g)$ in size polynomial in $|\mathrm{CNF}(g)|$.

Lemma 3.6. *Any sequence $(G_n : n \in \mathbb{N})$ of families G_n of n-ary Boolean functions of size $|G_n| \leq \exp\big(O(n \log n)\big)$ is encoding robust.*

Proof. Consider a Boolean function $g \in G_n$ and note that the CNF-encoding of g is of size $|\mathrm{CNF}(g)| \leq 2^n$. Let $\mathrm{CNF}(g) \coloneqq C_1 \wedge \cdots \wedge C_m$ and $F_g \coloneqq \left\{ f_g^{(i)} \coloneqq \bigwedge_{j \leq i} C_j \,\middle|\, i \in [m] \right\}$. Define $F_n \coloneqq \bigcup_{g \in G_n} F_g$ and observe that the tree-like proof system $\mathrm{sem}^\star(F_n : n \in \mathbb{N})$ can derive any $g \in G_n$ from $\mathrm{CNF}(g)$ in size $O\big(|\mathrm{CNF}(g)|\big)$ by iteratively deriving the functions $f_g^{(i)}$ for $i = 1, \ldots, m$. Since $|G_n| \leq \exp(O(n \log n))$ the family F_n is of size $|F_n| \leq \exp(O(n \log n))$ and the sequence $(G_n : n \in \mathbb{N})$ thus encoding robust. □

Let $G = (P, H, E)$ denote a bipartite graph with partitions of size $|P| = m + 1$ and $|H| = m$. The graph pigeonhole principle over G claims that the set of pigeons P fits into holes H such that each pigeon $p \in P$ flies to an adjacent hole $h \in H$ while each hole contains at most a single pigeon. Denote by $\mathrm{PHP}(G)$ the CNF encoding of the above principle over variables $\left\{ x_{\{p,h\}} \,\middle|\, \{p, h\} \in E \right\}$, each variable $x_{\{p,h\}}$ indicating that pigeon p flies to hole h, consisting of clauses

$$\bigvee_{h \in N(p)} x_{\{p,h\}} \qquad \forall p \in P \qquad \text{(pigeon axiom)}$$

$$\bar{x}_{\{p,h\}} \vee \bar{x}_{\{p',h\}} \qquad \forall h \in H, \forall p \neq p' \in N(h) \ . \qquad \text{(hole axiom)}$$

Consider the family of graph pigeonhole principles over n variables. We would like to argue that this family is simple. Since $\mathrm{sem}(\mathcal{F})$ is generally not closed under restrictions this is not immediate. Additionally, for $n \in \mathbb{N}$ there are $\exp\big(O(\binom{n^2}{n})\big)$ many graph pigeonhole principle formulas over n variables that do not contain an isolated pigeon and are therefore not immediately refuted. To argue that this family of formulas is simple we can thus not add a refutation for each instance to F_n. However, by introducing linear integer inequalities as proof lines this entire family can be refuted efficiently.

Proposition 3.7. *The family of graph pigeonhole principle formulas is simple.*

Proof. Consider the sequence $(F_n : n \in \mathbb{N})$ where each F_n is the set of threshold functions that can be either written as $\sum_{i \in I} x_i \geq j$ or $\sum_{i \in I} x_i \leq j$ for $I \subseteq [n]$ and $j \in [n]$. Since $|F_n| \leq \exp\big(O(n \log n)\big)$, by Lemma 3.6 these functions are encoding robust. Thus by an irrelevant increase in the size of each F_n we may derive the graph pigeonhole principle expressed with inequalities to then refute the formula. □

For an (ordinary) graph $G = (V, E)$ and a *charge vector* $\alpha \in \{0,1\}^V$ denote by $\mathrm{Tseitin}(G, \alpha)$ the CNF formula over variables $\{x_e \mid e \in E\}$ consisting of constraints

$$\sum_{e \ni v} x_e = \alpha_v \mod 2 \qquad \forall v \in V \tag{4}$$

each encoded as a CNF formula with $2^{\deg(v)-1}$ clauses. Consider the family of Tseitin formulas over n variables. This family is of size $\exp\big(O(\binom{n^2}{n})\big)$ if we ignore formulas that contain empty clauses and are thus trivial to refute. While this family is large it is easily seen that it can be efficiently refuted by proof systems with access to affine linear forms over $\mathbb{F}_2$.

Proposition 3.8. *The family of Tseitin contradictions is simple.*

Proof. Note that the sequence $\mathcal{F} := (F_n : n \in \mathbb{N})$ of $\mathbb{F}_2$-affine linear forms is small and thus encoding robust according to Lemma 3.6. Add the necessary proof lines to each F_n so that from the CNF encoding of the Tseitin contradiction $\mathrm{sem}^\star(\mathcal{F})$ can derive the corresponding system of $\mathbb{F}_2$-affine linear forms which may then be readily refuted. □

Finally, we consider lifted families of formulas. For a function $g : \{0,1\}^m \to \{0,1\}$ and a n-variate CNF formula A, let $A \circ g^n$ denote the lifted CNF formula where we replace every variable $x \in \mathrm{Vars}(A)$ by a formula computing g over a fresh set of variables, and writing it out as a CNF formula.

Proposition 3.9. *For any simple sequence $(\mathcal{A}_n : n \in \mathbb{N})$ of families $\mathcal{A}_n$ of n-variate CNF formulas and any gadget $g : \{0,1\}^m \to \{0,1\}$ it holds that the lifted sequence $\big(\{A \circ g^n \mid A \in \mathcal{A}_n\} : n \in \mathbb{N}\big)$ is also simple.*

Proof. Denote by G_{mn} the family of mn-variate Boolean functions obtained by viewing g as an mn-variate function depending on the first m variables and applying any permutation $\sigma : [mn] \to [mn]$ to g. Note that this family is not too large $|G_{mn}| \leq \exp\big(O(mn \log mn)\big)$ and thus encoding robust by Lemma 3.6. Denote by $(H_n : n \in \mathbb{N})$ the sequence certifying that $(G_n : n \in \mathbb{N})$ is encoding robust and let $(F_n : n \in \mathbb{N})$ be the sequence certifying that $(\mathcal{A}_n : n \in \mathbb{N})$ is simple. By composing each $f \in F_n$ with g we obtain the sequence $\big(H_{mn} \cup \{f \circ g^n \mid f \in F_n\} : n \in \mathbb{N}\big)$ with which the lifted formulas can be efficiently refuted. □

4 Lower Bounds for Tree-Like Cook–Reckhow Proof Systems

In this section we state our main result on tree-like semantic refutations. We then proceed to prove the corollaries mentioned in the introduction assuming our main theorem. The proof of the main theorem is deferred to Section 5.

To obtain the desired consequences mentioned in Section 1.3 we cannot simply measure a set of proof lines F_n by its size $|F_n|$. Instead, we have to resort to a more refined notion taking the structure of the clique formula into consideration. Recall that if a k-partite graph G with partition $V_1, \dots, V_k$ contains a k-clique, then this k-clique is a tuple $t \in \mathcal{T} \coloneqq \prod_{i \in [k]} V_i$. For any CNF formula $\mathrm{BlockClique}\colon \{0,1\}^X \times \{0,1\}^Y \to \{0,1\}$, as introduced in Section 2.4, there is by definition a one-to-one correspondence between tuples $t \in \mathcal{T}$ and assignments $\rho_t \in \{0,1\}^X$ such that for any graph $G \in \{0,1\}^Y$ it holds that $\mathrm{BlockClique}(\rho_t, G) = 1$ if and only if t is a k-clique in G. In what follows, instead of considering *all* assignments $\{0,1\}^X$, we only consider assignments $\{\rho_t \mid t \in \mathcal{T}\} \subseteq \{0,1\}^X$ corresponding to tuples. Since these assignments are in one-to-one correspondence with tuples we simply discuss tuples and it is understood that we really mean to discuss the corresponding assignments.

Associate each proof line $f \in F_{|X|}$ with the tuples it rules out: let $Q\colon F_{|X|} \to \{0,1\}^{\mathcal{T}}$ denote the map defined by $Q(f) \coloneqq \{t \in \mathcal{T} \mid f(\rho_t) = 1\}$ and extend this notation to sets $F \subseteq F_{|X|}$ by $Q(F) \coloneqq \{Q(f) \mid f \in F\}$. With this notation at hand we can state our main theorem.

Theorem 4.1. *For any CNF formula* $\mathrm{BlockClique}(X, Y)$ *as introduced in Section 2.4, any* $\varepsilon \in \mathbb{R}^+$*, and integer* $D, k, n \in \mathbb{N}$ *such that* $k < n^{\varepsilon/61}$ *and* $D \leq 2 \log n$ *the following holds. If* $\mathcal{F} \coloneqq (F_t \mid t \in \mathbb{N})$ *is a sequence of proof lines satisfying* $|Q(F_{|X|})| \leq \exp(n^{2-\varepsilon})$*, then for* $G \sim \mathcal{G}(n, k, n^{-2/D})$ *asymptotically almost surely any tree-like* $\mathrm{sem}(\mathcal{F})$ *refutation of* $\mathrm{BlockClique}(X, G)$ *is of length* $n^{\Omega(\varepsilon^2 D)}$*.*

Theorems 1.1 to 1.3 follow from Theorem 4.1 as explained next. Note that since $k \leq n$ the number of clauses in the different encodings of clique considered, i.e., $\mathrm{BlockClique}^{\mathrm{U}}(G, k)$, $\mathrm{BlockClique}_c(G, k)$, and $\mathrm{BlockClique}^{\mathrm{B}}(G, k)$, is bounded by $O(n^4)$. This is *independent* of the number of variables of the formulas and thus implies that for any of these encodings Theorem 4.1 yields a superpolynomial lower bound in the formula size, provided $D = \omega(1)$. For Theorem 1.1 let $k = 4 \log n$, $D = 2 \log n$, and consider the unary encoding $\mathrm{BlockClique}^{\mathrm{U}}(G, k)$ with the range axioms expressed as a conjunction of width 3 clauses by introducing $O(kn)$ extension variables. Similarly, Theorem 1.2 follows for the identical parameter setting with $\ell \coloneqq 2c$ but by considering $\mathrm{BlockClique}_c(G, k)$ with range axioms expressed as 3-CNF formulas by introducing extension variables. Finally, Theorem 1.3 follows from Theorem 4.1 for $D = \log^{\varepsilon_0} n$ and $k = 2 \log^{\varepsilon_0} n$ and noting that the number of variables $\mathrm{BlockClique}_c(G, k)$ is defined over ranges from $n \log^{\varepsilon_0} n$ to $2 \log^{1+\varepsilon_0} n$ as c ranges from 1 to $\lceil \log n \rceil$.

We defer the proof of Theorem 4.1 to Section 5, and first show how our results for concrete proof systems mentioned in Section 1.3 follow.

Let us first consider tree-like cutting planes (CP) refutations of $\mathrm{BlockClique}^{\mathrm{U}}(G, k)$. Each line of a CP refutation is a linear integer inequality, which may be viewed as an affine halfspace in $\mathbb{R}^{kn}$. Intuitively, since a halfspace is determined by kn points and we are only interested in the points $\rho_t \in \mathbb{R}^{kn}$ for tuples $t \in \mathcal{T}$ we see that there are approximately $\binom{n^k}{kn} \leq n^{k^2 n}$ many distinct proof lines. Since $k \ll \sqrt{n}$ we should be able to appeal to Theorem 4.1.

Instead of formalizing the above intuition into a rigorous argument it is more convenient to appeal to bounds on the VC-dimension of the family of affine halfspaces to then invoke Theorem 4.1 by appealing to Theorem 2.1.

Theorem 4.2 ([SSBD14])**.** *The* VC*-dimension of the family of affine halfspaces in* $\mathbb{R}^d$ *is* $d + 1$*.*

Our main result on tree-like cutting planes refutations follows.

Corollary 4.3. *For integer $D, k, n \in \mathbb{N}$ such that $k \leq n^{1/64}$ and $D \leq 2\log(n)$ it holds that tree-like cutting planes requires for $G \sim \mathcal{G}(n, k, n^{-2/D})$ asymptotically almost surely refutations of length $n^{\Omega(D)}$ to refute* $\mathrm{BlockClique}^{\mathrm{U}}(G, k)$.

Proof. Denote by $(F_\ell \mid \ell \in \mathbb{N})$ the sequence of families of linear threshold functions in ℓ dimensions. To bound the cardinality of $Q(F_{kn})$ view any linear threshold function as an affine halfspace partitioning the vertices of the kn-dimensional hypercube.

Note that if a subset of tuples $X \subseteq \mathcal{T}$ is shattered by the family $Q(F_{kn})$, then the corresponding set of assignments $\{\rho_t \mid t \in X\}$ is shattered by the family of affine halfspaces. By appealing to Theorem 4.2 we thus get that $\mathrm{VC}\big(Q(F_{kn})\big) \leq nk + 1$. The result follows by appealing to Theorem 4.1 with the unary k-clique formula $\mathrm{BlockClique}^{\mathrm{U}}(G, k)$, $\varepsilon := 61/64$, and the bound $|Q(F_{kn})| \leq (n^k)^{nk+1}$ obtained from Theorem 2.1. □

Note that Corollary 4.3 above can be slightly generalized to hold for inferential proof systems operating over low-width disjunctions of linear thresholds, i.e., bounded-width tree-like resolution over cutting planes, or equivalently bounded-depth *stabbing planes*. This follows from the fact that the set $F^{(s)}_{nk}$ of disjunctions of s linear threshold functions over nk Boolean variables satisfies $\big|Q\big(F^{(s)}_{nk}\big)\big| \leq \big|Q\big(F^{(1)}_{nk}\big)\big|^s \leq \exp\big(O(snk^2 \log n)\big)$. Hence as long as s is a small power of n we obtain optimal average-case clique refutation length lower bounds.

The second instantiation of Theorem 4.1 is for tree-like $\mathrm{Res}(\oplus)$.

Corollary 4.4. *For integer $D, k, n \in \mathbb{N}$ such that $k \leq n^{1/64}$ and $D \leq 2\log(n)$ it holds that tree-like* $\mathrm{Res}(\oplus)$ *requires for $G \sim \mathcal{G}(n, k, n^{-2/D})$ asymptotically almost surely refutations of length $n^{\Omega(D)}$ to refute* $\mathrm{BlockClique}^{\mathrm{U}}(G, k)$.

The proof of this result follows along the same lines as the proof of Corollary 4.3. Let us recall the bound on the VC-dimension of linear systems of equations over $\mathbb{F}_2$.

Theorem 4.5 ([CEYZ20])**.** *The* VC*-dimension of solutions of linear systems of equations over $\mathbb{F}_2$ in d variables is d.*

With this bound at hand the proof of Corollary 4.3 can be essentially repeated to obtain Corollary 4.4. For completeness we provide the proof next.

Proof of Corollary 4.4. Let F_{nk} denote the set of $\mathbb{F}_2$-affine linear forms over the kn variables of $\mathrm{BlockClique}^{\mathrm{U}}(G, k)$, and denote by $F^{\vee}_{kn}$ the set of F_{kn}-clauses, that is, $F^{\vee}_{kn} := \big\{\bigvee_{f \in I} f \,\big|\, I \subseteq F_{kn}\big\}$. It remains to bound the VC-dimension of $Q(F^{\vee}_{kn})$. Since every set of tuples shattered by $Q(F^{\vee}_{kn})$ is equivalently shattered by the set of solutions of linear systems of equations over $\mathbb{F}_2$ we may appeal to Theorem 4.5 to obtain that $\mathrm{VC}\big(Q(F^{\vee}_{kn})\big) \leq kn$. Theorem 4.1 thus yields the claimed bound if instantiated with the formula $\mathrm{BlockClique}^{\mathrm{U}}(G, k)$, $\varepsilon := 61/64$, and the bound $|Q(F^{\vee}_{kn})| \leq n^{k^2 n}$ obtained by appealing to Theorem 2.1 with $\mathrm{VC}\big(Q(F^{\vee}_{kn})\big) \leq kn$. □

The penultimate application of Theorem 4.1 is for tree-like threshold $\mathrm{Th}(d)$ proof systems, formally establishing Corollary 1.5.

Corollary 4.6. *For integer $D, k, n \in \mathbb{N}$ such that $k = \Theta(D)$ and $D \leq 2\log n$, it holds that tree-like* $\mathrm{Th}\big(\log n / 2\log\log n\big)$ *requires for $G \sim \mathcal{G}(n, k, n^{-2/D})$ asymptotically almost surely refutations of length $n^{\Omega(D)}$ to refute* $\mathrm{BlockClique}^{\mathrm{B}}(G, k)$.

Corollary 1.5 readily follows from Corollary 4.6 for $D = \log^{\varepsilon} n$, $k = 2\log^{\varepsilon} n$, and a re-parametrization by the number of variables $O(\log^{1+\varepsilon} n)$ the formula $\mathrm{BlockClique}^{\mathrm{B}}(G, 2\log^{\varepsilon} n)$ is defined over.

The proof of Corollary 4.6 follows the same template as the proofs of Corollaries 4.3 and 4.4. The following bound on the VC-dimension of polynomial threshold functions (PTFs) over Boolean variables can be deduced from [AB09] and appears explicitly in, e.g., [BFPJH21].

Theorem 4.7 ([AB09, BFPJH21])**.** *The* VC*-dimension of PTFs over t Boolean variables and degree at most d is $\sum_{i=0}^{\min(t,d)} \binom{t}{i}$.*

With this bound at hand the proof of Corollary 4.6 is easily established as follows.

Proof of Corollary 4.6. Let $k := \alpha D$ and $d := \log n / 2\log\log n$. Since $\mathrm{Th}(d)$ operates over PTFs of degree at most d by Theorem 4.7 it holds that the set $F_{k\log n}$ of PTFs over $k\log n$ variables satisfies

$$\mathrm{VC}\big(Q(F_{k\log n})\big) \leq (1 + k\log n)^d \leq (4\alpha \log n)^{2d} \leq n^{1+\mathrm{o}(1)} \quad . \tag{5}$$

We may thus appeal to Theorem 4.1 with $\mathrm{BlockClique}^{\mathrm{B}}(G, k)$ and the bound $|Q(F_{k\log n})| \leq n^{kn^{1+\mathrm{o}(1)}}$ obtained from Theorem 2.1. □

Last but not least we get our corollary for tree-like Frege over formulas of bounded size.

Corollary 4.8. *For any $\varepsilon \in \mathbb{R}^+$ and integer $c, D, k, n \in \mathbb{N}$ such that $D \leq 2\log n$, $0 < c \leq \lceil \log n \rceil$ and $k = O(D)$, it holds that tree-like Frege with line size $n^{2-\varepsilon}$ requires for $G \sim \mathcal{G}(n, k, n^{-2/D})$ asymptotically almost surely refutations of length $n^{\Omega(D)}$ to refute* $\mathrm{BlockClique}_c(G, k)$*.*

Corollary 1.6 follows from Corollary 4.8 for $D = \log^{\varepsilon_0} n$, $k = 2\log^{\varepsilon_0} n$, and observing that the number of variables $\mathrm{BlockClique}_c(G, k)$ is defined over ranges between $2\log^{1+\varepsilon_0} n$ and $n\log^{\varepsilon_0} n$ for $c \in [1, \lceil \log n \rceil]$.

Proof. The number of formulas of size s is bounded by $\exp\big(O(s\log s)\big)$. Hence the number of proof lines $F_{\lceil kcn^{1/c} \rceil}$ is bounded by $|F_{\lceil kcn^{1/c} \rceil}| \leq \exp(n^{2-\varepsilon}(2-\varepsilon)\log n) \leq \exp(n^{2-\varepsilon/2})$. We may thus appeal to Theorem 4.1 with the $\mathrm{BlockClique}_c(G, k)$ encoding and $\varepsilon/2$. □

5 Length Lower Bounds on Tree-Like Semantic Refutations

This section is devoted to arguing that for $G \sim \mathcal{G}(n, k, n^{-2/D})$ with $D \leq 2\log n$ and any family $\mathcal{Q} \subseteq 2^{\mathcal{T}}$ of size $|\mathcal{Q}| \leq \exp(n^{2-\delta})$ asymptotically almost surely there exists a *pseudo-measure* $\mu \colon 2^{\mathcal{T}} \to \mathbb{R}$, linear over the set of tuples $\mathcal{T}$, such that

1. $\mu(\mathcal{T}) \geq 1/2$, and
2. if $Q \in \mathcal{Q}$ is *ruled out by a missing edge*, that is, there is an edge $e \notin E(G)$ such that $e \subseteq t$ for all $t \in Q$, then $\mu(Q) \leq n^{-\Omega(D)}$.

Let us sketch the proof of Theorem 4.1 assuming there exists a pseudo-measure μ with the above properties.

Consider a tree-like $\mathrm{sem}(\mathcal{F})$ refutation $\pi := (f_1, \ldots, f_s)$ of a formula A. Towards contradiction let us assume that $s \leq n^{\lambda D}$ for some small enough $\lambda \in \mathbb{R}^+$. By standard arguments we can extract from π a balanced $F_{|X|}$-decision tree T solving $\mathrm{Search}(A)$ of size $|T| = O(s)$. Consider

a node $v \in V(T)$, denote the unique root-to-v path by $p_v := (u_1, u_2, \ldots, u_\tau = v)$, and let $C_v := \bigwedge_{u \in p_v}(f_u = b_u)$ be the conjunction of the answers to the queried functions f_u on p_v. Since T is balanced each C_v is a small conjunction of size $O(\log s)$. Hence the family

$$\{Q(C_v) \mid v \in V(T)\} \subseteq \mathcal{Q} := \left\{ \bigcap_{f \in I} Q(f) \cap \bigcap_{g \in J} Q(\neg g) \;\middle|\; I, J \subseteq F_{|X|},\ |I \sqcup J| \leq O(\log s) \right\} \tag{6}$$

is of size at most $|Q(F_{|X|})|^{O(\log s)}$. Since by assumption $|Q(F_{|X|})| \leq \exp(n^{2-\delta})$ we obtain that $|\mathcal{Q}| \leq \exp(n^{2-\delta/2})$ for small enough $\lambda > 0$.

We may thus consider the pseudo-measure μ for the family $\mathcal{Q}$. Associate each node $u \in V(T)$ with $\mu\big(Q(C_u)\big)$. Note that by linearity of the measure it holds that the values associated with the two out-neighbours v_0, v_1 of u sum to the value associated with u, that is, it holds that $\mu\big(Q(C_u)\big) = \mu\big(Q(C_{v_0})\big) + \mu\big(Q(C_{v_1})\big)$. By Property 1 of the measure the root node r of T is associated with $\mu\big(Q(C_r)\big) = \mu(\mathcal{T}) \geq 1/2$. If we can further argue that each leaf of the refutation is associated with a set $Q \in \mathcal{Q}$ ruled out by a missing edge, then we may conclude by Property 2 of μ that T has $n^{\Omega(D)}$ leaves and we thus obtain the claimed length lower bound on tree-like sem($\mathcal{F}$) refutations. This completes the proof sketch of Theorem 4.1 assuming the existence of a pseudo-measure μ with the above properties.

We rely on the same construction for the pseudo-measure μ as in [dRPR23]. The main difference lies in its analysis: whereas [dRPR23] could guarantee Property 2 for structured Q only, in our setting there is no restriction on the nature of Q—each Q is an arbitrary set of tuples that are all ruled out by the same missing edge. The only guarantee we have is that there are not too many such sets; that is, that $\mathcal{Q}$ is somewhat small. We show that the analysis of [dRPR23] is flexible enough to be adapted to our setting.

Organization. In Section 5.1 we formalize the above proof outline. We define the notion of *cores* from [dRPR23] in Section 5.2 and define the property of random graphs we need for the lower bound to hold. In Section 5.3 we establish the required properties of the pseudo-measure and end with Section 5.4 arguing that random graphs are pseudorandom.

5.1 The Pseudo-Measure and Proof of the Main Theorem

Denote by vc(E) the *minimum vertex cover* of a set of edges E, let $G \sim \mathcal{G}(n, k, p)$ be a k-partite graph, and denote by e a *potential edge*, that is, an edge that has non-zero probability of being sampled by $\mathcal{G}(n, k, p)$. Consider the biased characters

$$\chi_e(G) := \begin{cases} \frac{1-p}{p} & \text{if } e \in E(G) \\ -1 & \text{otherwise,} \end{cases} \tag{7}$$

and for a set of potential edges E define

$$\chi_E(G) := \prod_{e \in E} \chi_e(G) \ . \tag{8}$$

For $d \in \mathbb{N}$ the *pseudo-measure* $\mu_{G,d} \colon 2^{\mathcal{T}} \to \mathbb{R}$ is defined on tuples $t \in \mathcal{T}$ by

$$\mu_{G,d}(t) := n^{-k} \sum_{\substack{E \subseteq \binom{t}{2}: \\ \mathrm{vc}(E) \leq d}} \chi_E(G) \ , \tag{9}$$

and extended linearly to sets of tuples $Q \subseteq \mathcal{T}$

$$\mu_{G,d}(Q) \coloneqq \sum_{t \in Q} \mu_{G,d}(Q) \ . \tag{10}$$

Next we state the properties required of the pseudo-measure to deduce Theorem 4.1.

Theorem 5.1. *There is a constant $c \in \mathbb{R}^+$ such that for small enough $\varepsilon \in \mathbb{R}^+$ and for all $d, D, \delta \in \mathbb{R}^+$ and integer $k, n \in \mathbb{N}$ satisfying $\delta > c\varepsilon$, $k < n^{\delta/60}$, $D \leq 2\log n$, and $d = \varepsilon D$ the following holds for any family $\mathcal{Q} \subseteq 2^{\mathcal{T}}$ of size $|\mathcal{Q}| \leq \exp(n^{2-\delta})$. If $G \sim \mathcal{G}(n, k, n^{-2/D})$, then asymptotically almost surely*

1. *$\mu_{G,d}$ is linear over $\mathcal{T}$,*
2. *$\mu_{G,d}(\mathcal{T}) = 1 - o(1)$, and*
3. *for all $Q \in \mathcal{Q}$ ruled out by a missing edge it holds that $\mu_{G,d}(Q) \leq n^{-\Omega(\varepsilon d)}$.*

Note that the first property of Theorem 5.1 is immediate by the definition of the pseudo-measure. The second property follows from the analysis of [dRPR23]. All that remains is to argue the third property. Before proving Theorem 5.1 let us show that Theorem 4.1 follows from Theorem 5.1. For convenience we restate Theorem 4.1.

Theorem 4.1. *For any CNF formula* $\mathrm{BlockClique}(X, Y)$ *as introduced in Section 2.4, any $\varepsilon \in \mathbb{R}^+$, and integer $D, k, n \in \mathbb{N}$ such that $k < n^{\varepsilon/61}$ and $D \leq 2\log n$ the following holds. If $\mathcal{F} \coloneqq (F_t \mid t \in \mathbb{N})$ is a sequence of proof lines satisfying $|Q(F_{|X|})| \leq \exp(n^{2-\varepsilon})$, then for $G \sim \mathcal{G}(n, k, n^{-2/D})$ asymptotically almost surely any tree-like* $\mathrm{sem}(\mathcal{F})$ *refutation of* $\mathrm{BlockClique}(X, G)$ *is of length $n^{\Omega(\varepsilon^2 D)}$.*

Proof. Let π be a $\mathrm{sem}^\star(\mathcal{F})$ refutation of $\mathrm{BlockClique}_c(G, k)$ and suppose that $|\pi| \leq n^{\lambda D}$ for a small enough $\lambda \coloneqq \lambda(\varepsilon) \in \mathbb{R}^+$. The first step in the proof is to construct an $O\big(\log|\pi|\big)$-depth $F_{|X|}$-decision tree T that solves the falsified clause search problem $\mathrm{Search}\big(\mathrm{BlockClique}(X, G)\big)$. In a second step we then argue that the tree T must have at least $n^{\Omega(\varepsilon d)}$ leaves, thereby establishing the claimed length lower bound on the refutation π.

Since the proof π is tree-like, we can construct T inductively via a Brent-Spira balancing argument: find a line $f \in \pi$ such that the subtree π_f of π rooted at f is of size $|\pi|/3 < |\pi_f| < 2|\pi|/3$. Query f and consider two cases. If $f = 0$, then one of the leaves of π_f is a solution to the falsified clause search problem, since the leaves of π_f imply f. Otherwise, if $f = 1$, then the pruned proof tree $\pi \setminus \pi_f$ will contain an axiom falsified by any assignment agreeing with the query. In both cases, we can recurse on a tree-like proof of size at most $2|\pi|/3$. By induction we thus obtain a tree T solving the falsified clause search problem in depth $\mathrm{depth}(T) = O\big(\log|\pi|\big) \leq \Delta \coloneqq c\lambda D \log n$ for a large enough constant $c \in \mathbb{R}^+$. This establishes the first step of the argument. It remains to argue that T has many leaves.

Associate with each leaf $\ell \in T$ the set of tuples Q_ℓ containing $t \in Q_\ell$ whose corresponding assignment ρ_t is consistent with the set of queries on the root-to-ℓ path. Every such set Q_ℓ is the intersection of at most $\mathrm{depth}(T) \leq \Delta$ sets of the form $Q(f)$ or $Q(\neg f)$ for $f \in F_{|X|}$. Consider the set

$$\mathcal{Q} \coloneqq \left\{ \bigcap_{f \in I} Q(f) \cap \bigcap_{g \in J} Q(\neg g) \;\middle|\; I, J \subseteq F_{|X|},\ |I \sqcup J| \leq \Delta \right\} \ . \tag{11}$$

It holds that $\{Q_\ell \mid \ell \in \mathrm{leaves}(T)\} \subseteq \mathcal{Q}$ and $|\mathcal{Q}| \leq \big(2|Q(F_{|X|})|\big)^\Delta < 2^{\Delta(n^{2-\varepsilon}+1)}$ by the assumption on the size of $Q\big(F_{|X|}\big)$. Since $\Delta = c\lambda D \log n$ and $D \leq 2\log n$, for large enough n and $\delta \coloneqq 60\varepsilon/61$

it thus holds that $|Q| \leq \exp(n^{2-\delta})$ and $k < n^{\delta/60}$. This allows us to consider the pseudo-measure $\mu := \mu_{G,\varepsilon_0 D}$ as in Theorem 5.1 for the set Q and $\varepsilon_0 := \delta/c_0 = \varepsilon \cdot 60/61 c_0$ for a large enough constant $c_0 \in \mathbb{R}^+$.

Because T solves the falsified clause search problem there is a clause $C \in \text{BlockClique}(X, G)$ falsified by all tuples $t \in Q_\ell$. In other words, for all $t \in Q_\ell$ it holds that the assignment ρ_t falsifies C. Fix a leaf $\ell \in \text{leaves}(T)$, a tuple $t \in Q_\ell$, and denote by C the associated falsified clause. Recall that the formula $\text{BlockClique}(X, G)$ is obtained by restricting the $\text{BlockClique}(X, Y)$ formula by $G \in \{0,1\}^Y$, where each variable $y_e \in Y$ is the indicator variable of a potential edge e.

Consider the formula $F_C \subseteq \text{BlockClique}(X, Y)$ consisting of all clauses $C' \supseteq C$. We claim that every clause $C' \in F_C$ contains a positive literal y_e with $e \subseteq t$: recall that every clause contains at most one literal over Y and consider the graph $G_t : Y \mapsto \{0,1\}$ with edges $y_e \mapsto 1$ if and only if $e \subseteq t$. Since $\text{BlockClique}(\rho_t, G_t) = 1$ but $C\restriction_{\rho_t} = 0$, it holds that every clause $C' \in F_C$ either contains a positive literal y_e for $e \subseteq t$, or a negative literal $\bar{y}_e$ for an edge $e \not\subseteq t$. Suppose some clause C' contains a negative literal $\bar{y}_e$ for $e \not\subseteq t$. Since $C'\restriction_{y_e \mapsto 1} = C$ which maps to false under ρ_t we see that $\text{BlockClique}(\rho_t, G_t \vee \{y_e \mapsto 1\}) = 0$. But this cannot be since t is a clique in $G_t \vee \{y_e \mapsto 1\}$ and by definition ρ_t is thus a satisfying assignment of $\text{BlockClique}(X, G_t \vee \{y_e \mapsto 1\})$.

We conclude that $F_C = \left(\bigwedge_{e \in E} y_e\right) \vee C$ for some non-empty edge set $E \subseteq \binom{t}{2}$. Since the above holds for every $t \in Q_\ell$ it holds that $E \subseteq \bigcap_{t \in Q_\ell} \binom{t}{2}$. At least one edge $e \in E$ is not present in G as otherwise $C \notin \text{BlockClique}(X, G)$. Thus by the third property of the pseudo-measure

$$\mu_{G,\varepsilon_0 D}(Q_\ell) = n^{-\Omega(\varepsilon_0^2 D)} = n^{-\Omega(\varepsilon^2 D)} \tag{12}$$

for all leaves $\ell \in \text{leaves}(T)$. Since the sets $\{Q_\ell \mid \ell \in \text{leaves}(T)\}$ partition the set of tuples $\mathcal{T}$ and the measure is linear over $\mathcal{T}$ it holds that

$$\sum_{\ell \in \text{leaves}(T)} \mu(Q_\ell) = \mu(\mathcal{T}) \ . \tag{13}$$

By the first property we have that $\mu(\mathcal{T}) = 1 - o(1)$ and T thus has $n^{\Omega(\varepsilon^2 D)}$ many leaves; it is of depth $\text{depth}(T) = \Omega\big(\varepsilon^2 D \log(n)\big)$. Since $\text{depth}(T) = O\big(\log|\pi|\big)$ this contradicts the original assumption $|\pi| \leq n^{\lambda D}$. This completes the proof of Theorem 4.1. □

5.2 On Cores and Pseudorandom Graphs

This section recalls the notion of a core crucial to the analysis as originally defined in [dRPR23] to then define our notion of pseudorandom graphs.

Consider a graph H with vertex set $[k]$. In what follows we often identify H as a graph over the vertex set $[k]$ and at the same time treat it as a set of edges $H \subseteq \binom{[k]}{2}$. Slightly non-standard, let us say that a graph F is a *vertex-induced subgraph* of H if $V(F) = V(H) = [k]$ and there is a set $S \subseteq [k]$ such that $e \in E(F)$ if and only if $e \in E(H)$ and $e \subseteq S$. A *core* of a graph H is a vertex-induced subgraph F of H such that any minimum vertex cover of F is also a vertex cover of H.

Lemma 5.2 ([dRPR24, Theorem 4.4])**.** *There is a map* core *that maps graphs to one of its cores with the following property. For every graph F in the image of* core *it holds that $|V(E(F))| \leq 3\,\text{vc}(F)$ and that there exists an edge set $E_F^\star \subseteq V\big(E(F)\big) \times \big([k] \setminus V\big(E(F)\big)\big)$ such that $\text{core}(H) = F$ if and only if $E(H) = E(F) \sqcup E$ for $E \subseteq E_F^\star$.*

Going forward we fix the mapping core as exhibited in Lemma 5.2 and refer to $\mathrm{core}(H)$ as *the* core of H. Let us say that a graph H is in the *e-boundary* for an edge $e \in \binom{V(H)}{2}$ if and only if $\mathrm{vc}(H \cup e) > \mathrm{vc}(H)$. The notions of an e-boundary and a core interact nicely as stated next.

Proposition 5.3 ([dRPR24, Proposition 4.3])**.** *A core of a graph H is in the e-boundary if and only if H is.*

In other words Proposition 5.3 guarantees that either the entire family $\mathcal{H}_F := \{F \sqcup E \mid E \subseteq E_F^\star\}$ is in the e-boundary or no graph $H \in \mathcal{H}_F$ is in the e-boundary. The next statement gives a simple bound on the number of core graphs $\mathcal{F}_d := \left\{\mathrm{core}(H) \middle| H \subseteq \binom{[k]}{2} \text{ with } \mathrm{vc}(H) \leq d\right\}$ for graphs over k vertices and minimum vertex cover bounded by d.

Lemma 5.4 ([dRPR24, Lemma 2.4])**.** *There are at most $2^{b(a+\log k)}$ graphs H over k vertices with a vertex cover of size a and $|V(E(H))| \leq b$.*

For a tuple $t \in \mathcal{T}$ write $t_i := t \cap V_i$ for the vertex in the ith block V_i, consider a graph $H \subseteq \binom{[k]}{2}$, and denote by $H(t)$ the graph obtained by replacing each vertex $i \in [k]$ of H by t_i. In other words, the graph $H(t)$ is defined over vertices $\{t_1, \ldots, t_k\}$ and edges $\left\{\{t_i, t_j\} \middle| \{i,j\} \in E(H)\right\}$.

Definition 5.5 ($(\mathcal{Q}, D, \delta)$-good)**.** Let $s \in \mathbb{N}^+$, denote by $Q \subseteq \mathcal{T}$ a set of tuples, and consider a core F. A k-partite graph G with partition $V_1, \ldots, V_k$ of size $|V_i| = n$ each is *s-bounded over Q and F* if it holds that

$$n^{-k} \left| \sum_{t \in Q} \sum_{H \in \mathcal{H}_F} \chi_{H(t)}(G) \right| \leq s \ .$$

For $D \in \mathbb{R}^+$ and for a family $\mathcal{Q} \subseteq 2^{\mathcal{T}}$ of sets of tuples the graph G is *$(\mathcal{Q}, D, \delta)$-good* if G is s-bounded over Q and F for all $Q \in \mathcal{Q}$, all non-empty $F \in \mathcal{F}_{D/4}$, and $s := n^{2|E(F)|/D - \delta\,\mathrm{vc}(F)/10}$.

Theorem 5.6. *The following holds for $k, n \in \mathbb{N}$ and $D, \delta \in \mathbb{R}^+$ satisfying $D \leq 2\log n$ and $k \leq n^{1/5}$. For a family $\mathcal{Q} \subseteq 2^{\mathcal{T}}$ of size $|\mathcal{Q}| \leq \exp(n^{2-\delta})$ it holds that $G \sim \mathcal{G}(n, k, n^{-2/D})$ is asymptotically almost surely $(\mathcal{Q}, D, \delta)$-good.*

We defer the proof of Theorem 5.6 to Section 5.4. In the following we verify that Theorem 5.6 suffices to prove Theorem 5.1.

5.3 The Pseudo-Measure Is Well-Behaved on Good Graphs

As mentioned earlier it was already shown in [dRPR23] that the measure is large on the entire set of tuples $\mathcal{T}$. Since we are using a different notion of pseudorandomness we reprove it for completeness.

Lemma 5.7. *There is a constant $c \in \mathbb{R}^+$ such that for small enough $\varepsilon \in \mathbb{R}^+$ and for all $d, D, \delta \in \mathbb{R}^+$ and integer $k, n \in \mathbb{N}$ satisfying $\delta > c\varepsilon$, $k \leq n^{\delta/60}$, $D \leq 2\log n$, and $d = \varepsilon D$ the following holds for any family $\mathcal{Q} \subseteq 2^{\mathcal{T}}$ of size $|\mathcal{Q}| \leq \exp(n^{2-\delta})$. If $\mathcal{T} \in \mathcal{Q}$ and G is $(\mathcal{Q}, D, \delta)$-good, then $\mu_{G,d}(\mathcal{T}) \geq 1 - n^{-\Omega(1)}$.*

Proof. By the triangle inequality and Lemma 5.2 we may write

$$\mu_{G,d}(\mathcal{T}) = n^{-k} \sum_{t\in\mathcal{T}} \sum_{\substack{H\subseteq\binom{[k]}{2}:\\ \mathrm{vc}(H)\le d}} \chi_{H(t)}(G) \tag{14}$$

$$= n^{-k} \sum_{t\in\mathcal{T}} \chi_{\emptyset(t)}(G) + n^{-k} \sum_{\substack{F\in\mathcal{F}_d:\\ F\neq\emptyset}} \sum_{t\in\mathcal{T}} \sum_{H\in\mathcal{H}_F} \chi_{H(t)}(G) \tag{15}$$

$$\ge 1 - \sum_{\substack{F\in\mathcal{F}_d:\\ F\neq\emptyset}} n^{-k} \Big| \sum_{t\in\mathcal{T}} \sum_{H\in\mathcal{H}_F} \chi_{H(t)}(G) \Big| \tag{16}$$

$$\ge 1 - \sum_{\substack{F\in\mathcal{F}_d:\\ F\neq\emptyset}} n^{2|E(F)|/D - \delta\,\mathrm{vc}(F)/10} \,, \tag{17}$$

where for the final inequality we relied on G being $(\mathcal{Q}, D, \delta)$-good and the fact that $\mathcal{T} \in \mathcal{Q}$. Since according to Lemma 5.2 cores $F \in \mathcal{F}_d$ have few non-isolated vertices $|V(E(F))| \le 3\,\mathrm{vc}(F)$ it holds that $|E(F)| \le 3d\,\mathrm{vc}(F)$, where we used that $\mathrm{vc}(F) \le d$. Summing over $i = \mathrm{vc}(F)$ and combining (17) with the above bound and the bound on the number of cores from Lemma 5.4 we obtain

$$\mu_{G,d}(\mathcal{T}) \ge 1 - \sum_{i\in[d]} 2^{3i(d+\log k)} n^{-i(\delta/10 - 6d/D)} \tag{18}$$

$$\ge 1 - \sum_{i\in[d]} n^{-i(\delta/10 - 12\varepsilon - 3\log k/\log n)} \,, \tag{19}$$

using that $i \le d = \varepsilon D \le 2\varepsilon \log n$. Since $k \le n^{\delta/60}$ and we may assume that $\delta > 240\varepsilon$ we conclude that $\mu_{G,d}(\mathcal{T}) \ge 1 - n^{-\Omega(1)}$. □

To prove Theorem 5.1 it thus remains to establish that the pseudo-measure is small on sets of tuples $Q \subseteq \mathcal{T}$ ruled out by a missing edge as summarized in the following statement.

Lemma 5.8. *There is a constant $c \in \mathbb{R}^+$ such that for small enough $\varepsilon \in \mathbb{R}^+$ and for all $d, D, \delta \in \mathbb{R}^+$ and integer $k, n \in \mathbb{N}$ satisfying $\delta > c\varepsilon$, $k < n^{\delta/60}$, $D \le 2\log n$, and $d = \varepsilon D$ the following holds for any family $\mathcal{Q} \subseteq 2^{\mathcal{T}}$ of size $|\mathcal{Q}| \le \exp(n^{2-\delta})$. If a graph G is $(\mathcal{Q}, D, \delta)$-good, then any $Q \in \mathcal{Q}$ ruled out by a missing edge satisfies $\mu_{G,d}(Q) \le n^{-\Omega(\varepsilon d)}$.*

Note that Theorem 5.1 follows from Lemmas 5.7 and 5.8 in combination with Theorem 5.6.

Proof. Fix $Q \in \mathcal{Q}$ ruled out by a missing edge, let e denote the corresponding edge, and denote by $i \neq j$ the indices of the blocks of the endpoints of e. Observe that for a tuple $t \in \mathcal{T}$ and a graph $H \subseteq \binom{[k]}{2}$ the identity $\chi_{H(t)\cup e} + \chi_{H(t)} = 0$ holds whenever $e \notin E(G)$ and $e \notin H(t)$. This identity allows us to pair a graph H that does not contain the edge $\{i, j\}$ with the graph $H \cup \{i, j\}$ to obtain that

$$\mu_{G,d}(Q) = n^{-k} \sum_{t\in Q} \sum_{\substack{H:\\ \mathrm{vc}(H)\le d}} \chi_{H(t)}(G) \tag{20}$$

$$= n^{-k} \sum_{t\in Q} \Big(\sum_{\substack{H:\, e\in H(t)\\ \mathrm{vc}(H)\le d}} \chi_{H(t)}(G) + \sum_{\substack{H:\, e\notin H(t)\\ \mathrm{vc}(H)\le d}} \chi_{H(t)}(G) \Big) \tag{21}$$

$$= n^{-k} \sum_{t\in Q} \sum_{\substack{H:\ \mathrm{vc}(H)=d\\ \mathrm{vc}(H\cup\{i,j\})>d}} \chi_{H(t)}(G) \,. \tag{22}$$

Note that the set of graphs we sum over are precisely the graphs in the $\{i,j\}$-boundary with vertex cover d. By Proposition 5.3 the families $\mathcal{H}_F$ for $F \in \mathcal{F}_d$ in the $\{i,j\}$-boundary and vertex cover $\mathrm{vc}(F) = d$ partition this set of graphs We thus obtain that

$$\mu_{G,d}(Q) = n^{-k} \sum_{\substack{F \in \mathcal{F}_d : \\ \mathrm{vc}(F)=d \\ \mathrm{vc}(F \cup \{i,j\})>d}} \sum_{t \in Q} \sum_{H \in \mathcal{H}_F} \chi_{H(t)}(G) \tag{23}$$

$$\leq \sum_{\substack{F \in \mathcal{F}_d : \ \mathrm{vc}(F)=d \\ \mathrm{vc}(F \cup \{i,j\})>d}} n^{2|E(F)|/D - \delta d/10} \ , \tag{24}$$

where we relied on the assumption that G is bounded over Q and F. Since according to Lemma 5.2 it holds that cores $F \in \mathcal{F}_d$ have few non-isolated vertices $|V(E(F))| \leq 3\,\mathrm{vc}(F)$ it also holds that $|E(F)| \leq 3d^2$, further using that $\mathrm{vc}(F) = d$. Appealing to Lemma 5.4 to bound the sum in (24), and using the above bound and the fact that $d = \varepsilon D \leq 2\varepsilon \log n$ we get that

$$\mu_{G,d}(Q) \leq 2^{3d(d+\log k)} \cdot n^{-d(\delta/10 - 6\varepsilon)} \tag{25}$$

$$\leq n^{-d(\delta/10 - 12\varepsilon - 3\log k/\log n)} \tag{26}$$

$$= n^{-\Omega(\varepsilon^2 D)} \ , \tag{27}$$

assuming that $k \leq n^{\delta/60}$, and that $\delta > 240\varepsilon$. This concludes the proof of Lemma 5.8. □

5.4 Random Graphs are Good

This section is devoted to the proof of Theorem 5.6. To this end we need to recall two further statements from [dRPR23]. If we call a vertex set t an *a-tuple* if $|t| = a$ and there is a tuple $t' \in \mathcal{T}$ such that $t \subseteq t'$, then the first statement claims that for any block V_i the size of the common neighbourhood $N^\cap(t, V_i) \coloneqq V_i \cap \bigcap_{u \in t} N(u)$ in V_i for graphs $G \sim \mathcal{G}(n,k,p)$ is tightly concentrated around its expected value $\mathbb{E}[|N^\cap(t,V_i)|] = p^{|t|}n$.

Lemma 5.9 ([dRPR24, Lemma 8.1]). *For any integer $D, k, n \in \mathbb{N}$ satisfying $k \leq n^{1/5}$ and $D \leq 2\log n$ the following holds for $G \sim \mathcal{G}(n,k,n^{-2/D})$ asymptotically almost surely. For any $i \in [k]$, any $a \leq D/4$, and any a-tuple $t \subseteq V(G) \setminus V_i$ it holds that $|N^\cap(t,V_i)| \in (1 \pm 1/k)\,\mathbb{E}[|N^\cap(t,V_i)|]$.*

The proof of Lemma 5.9 follows by the usual Chernoff bound plus union bound argument. To argue Theorem 5.6 we need bounds on linear combinations of Fourier characters associated with the same core F as stated next. For a set $A \subseteq [k]$ and a tuple $t \in \mathcal{T}$ let $t_A \coloneqq \{t_i \mid i \in A\}$, and extend this notation to $\mathcal{T}$ in the natural way $\mathcal{T}_A \coloneqq \{t_A \mid t \in \mathcal{T}\}$.

Lemma 5.10 ([dRPR24, Lemma 8.5]). *Consider a non-empty graph F over the vertex set $[k]$, let $A \coloneqq V\big(E(F)\big)$, and denote by $M \subseteq E(F)$ a matching in F. For any even $m \leq n^2$, any $r \in \mathbb{R}^+$, and any function $\xi : \mathcal{T}_A \to [-r,r]$ it holds that*

$$\Pr_{G[V_A]}\left[\left|\sum_{t \in \mathcal{T}_A} \chi_{F(t)}(G)\xi(t)\right| > s\right] \leq \left(\frac{r \cdot p^{-|E(F)|} \cdot \big(m/n^2\big)^{|M|/2} \cdot n^{|A|}}{s}\right)^m \ ,$$

where each edge $e \in \binom{V_A}{2}$ is sampled independently with probability p, provided e has endpoints in distinct blocks.

Lemma 5.10 is in fact a weaker statement than Lemma 8.5 from [dRPR24] since it fixes $Q := \mathcal{T}_A$. From Lemmas 5.9 and 5.10 we may derive Theorem 5.6 restated here for convenience.

Theorem 5.6. *The following holds for $k, n \in \mathbb{N}$ and $D, \delta \in \mathbb{R}^+$ satisfying $D \leq 2\log n$ and $k \leq n^{1/5}$. For a family $\mathcal{Q} \subseteq 2^{\mathcal{T}}$ of size $|\mathcal{Q}| \leq \exp(n^{2-\delta})$ it holds that $G \sim \mathcal{G}(n, k, n^{-2/D})$ is asymptotically almost surely $(\mathcal{Q}, D, \delta)$-good.*

The remainder of this section is devoted to the proof of Theorem 5.6. Let $p := n^{-2/D}$, fix a core F, let $A := V(E(F))$, and let $B := [k] \setminus A$. For the following it is convenient to write a tuple $t \in \mathcal{T}$ as $t = (t_A, t_B)$. For a graph G denote by $G_A := G[V_A]$ the graph induced by vertices $V_A := \bigcup_{i \in A} V_i$ and let $G_{\bar{A}} := \left(V_{[k]}, E(G) \setminus \binom{V_A}{2}\right)$ denote the remainder. For a fixed set of tuples $Q \in \mathcal{Q}$ and $G_{\bar{A}}$ fixed define

$$\xi_{F,Q,G_{\bar{A}}}(t_A) := p^{-|E_F^\star|} \sum_{t_B \in \mathcal{T}_B} \mathbb{1}_{\{(t_A,t_B)\in Q\}}\, \mathbb{1}_{\{E_F^\star(t_A,t_B) \text{ present}\}}(G_{\bar{A}}) \ . \tag{28}$$

Recall further that $\mathcal{H}_F = \{F \sqcup E \mid E \supseteq E_F^\star\}$. This allows us to write the considered sum as

$$\begin{aligned}
\sum_{t\in Q} \sum_{H \in \mathcal{H}_F} \chi_{H(t)}(G) &= \sum_{t \in Q} \chi_{F(t)}(G) \sum_{E \subseteq E_F^\star} \chi_{E(t)}(G) && (29)\\
&= \sum_{t \in Q} \chi_{F(t)}(G_A)\, p^{-|E_F^\star|}\, \mathbb{1}_{\{E_F^\star(t) \text{ present}\}}(G_{\bar{A}}) && (30)\\
&= \sum_{t_A \in \mathcal{T}_A} \chi_{F(t_A)}(G_A)\, p^{-|E_F^\star|} \sum_{t_B \in \mathcal{T}_B} \mathbb{1}_{\{(t_A,t_B)\in Q\}}\, \mathbb{1}_{\{E_F^\star(t_A,t_B) \text{ present}\}}(G_{\bar{A}}) && (31)\\
&= \sum_{t_A \in \mathcal{T}_A} \chi_{F(t_A)}(G_A)\, \xi_{F,Q,G_{\bar{A}}}(t_A) \ , && (32)
\end{aligned}$$

where we observe that since every edge in $E_F^\star$ has an endpoint outside A, the function $\xi_{F,Q,G_{\bar{A}}}$ only depends on the edges of $G_{\bar{A}}$ and not on those of G_A.

For $r \in \mathbb{R}^+$ denote by $\boldsymbol{X}_r(F, Q, G)$ the indicator random variable of the event $|\xi_{F,Q,G_{\bar{A}}}| \leq r$, and for $s \in \mathbb{R}^+$ let $\boldsymbol{Y}_s(F, Q, G)$ be the indicator random variable for the event

$$\left| \sum_{t\in Q} \sum_{H \in \mathcal{H}_F} \chi_{H(t)}(G) \right| \leq s \ . \tag{33}$$

Note that it holds that

$$\begin{aligned}
&\Pr_G\left[\exists F \in \mathcal{F}_{D/4}, Q \in \mathcal{Q} \text{ such that } \neg \boldsymbol{Y}_s(F, Q, G)\right]\\
&\quad = \Pr_G\left[\exists F \in \mathcal{F}_{D/4}, Q \in \mathcal{Q} \text{ such that } \neg \boldsymbol{Y}_s(F, Q, G) \text{ and } \neg \boldsymbol{X}_r(F, Q, G)\right] + && (34)\\
&\qquad \Pr_G\left[\exists F \in \mathcal{F}_{D/4}, Q \in \mathcal{Q} \text{ such that } \neg \boldsymbol{Y}_s(F, Q, G) \text{ and } \boldsymbol{X}_r(F, Q, G)\right]\\
&\quad \leq \Pr_G\left[\exists F \in \mathcal{F}_{D/4}, Q \in \mathcal{Q} \text{ such that } \neg \boldsymbol{X}_r(F, Q, G)\right] +\\
&\qquad \sum_{\substack{F \in \mathcal{F}_{D/4}\\ Q \in \mathcal{Q}}} \Pr_G\left[\neg \boldsymbol{Y}_s(F, Q, G) \,\middle|\, \boldsymbol{X}_r(F, Q, G)\right] \ . && (35)
\end{aligned}$$

To argue that $G \sim \mathcal{G}(n, k, p)$ is asymptotically almost surely $(\mathcal{Q}, D, \delta)$-good, we need to show that the initial probability in (34) is bounded by o(1) for $s := n^k p^{-|E(F)|} n^{-\delta\, \mathrm{vc}(F)/10}$. Theorem 5.6 thus follows from the following two claims.

Claim 5.11. *For $r := 3n^{k-|A|}$ it holds that*

$$\Pr_G\left[\exists F \in \mathcal{F}_{D/4}, Q \in \mathcal{Q} \text{ such that } \neg \boldsymbol{X}_r(F,Q,G)\right] = o(1) \ .$$

Claim 5.12. *For $r := 3n^{k-|A|}$ and $s := 6n^k p^{-|E(F)|}(4n^{\delta/2})^{-\operatorname{vc}(F)/4}$ it holds that*

$$\sum_{\substack{F\in\mathcal{F}_{D/4}\\ Q\in\mathcal{Q}}} \Pr_G\left[\neg \boldsymbol{Y}_s(F,Q,G) \,\middle|\, \boldsymbol{X}_r(F,Q,G)\right] = o(1) \ .$$

Proof of Claim 5.11. The probability that any such event fails can be bounded by Lemma 5.9 as follows. Lemma 5.9 implies that asymptotically almost surely for any core $F \in \mathcal{F}_{D/4}$, any set of tuples $Q \in \mathcal{Q}$, any set $A \subseteq [k]$ of size $|A| \le D/4$, and any $t_A \in Q_A$ it holds that

$$\begin{aligned}
\xi_{F,Q,G_{\bar{A}}}(t_A) &= p^{-|E_F^\star|} \sum_{t_B\in\mathcal{T}_B} \mathbb{1}_{\{(t_A,t_B)\in Q\}}\, \mathbb{1}_{\{E_F^\star(t_A,t_B) \text{ present}\}}(G_{\bar{A}}) && (36)\\
&\le p^{-|E_F^\star|} \sum_{t_B\in\mathcal{T}_B} \mathbb{1}_{\{E_F^\star(t_A,t_B) \text{ present}\}}(G_{\bar{A}}) && (37)\\
&\le p^{-|E_F^\star|} n^{k-|A|} (1+1/k)^{k-|A|}\, p^{|E_F^\star|} && (38)\\
&\le 3n^{k-|A|} = r \ . && \square
\end{aligned}$$

Proof of Claim 5.12. Fix a core $F \in \mathcal{F}_{D/4}$ and a set $Q \in \mathcal{Q}$. We claim that Lemma 5.10 implies that

$$\begin{aligned}
\Pr_G\left[\neg \boldsymbol{Y}_s(F,Q,G) \,\middle|\, \boldsymbol{X}_r(F,Q,G)\right] &= \Pr_G\left[\left|\sum_{t_A\in\mathcal{T}_A} \chi_{F(t_A)}(G)\, \xi_{F,Q,G_{\bar{A}}}(t_A)\right| > s \,\middle|\, \boldsymbol{X}_r(F,Q,G)\right] && (39)\\
&\le \left(\frac{r\cdot p^{-|E(F)|}\cdot (m/n^2)^{\operatorname{vc}(F)/4}\cdot n^{|A|}}{s}\right)^m && (40)
\end{aligned}$$

for any $s \in \mathbb{R}^+$ and any even integer $m \in \mathbb{N}$. Indeed, since any graph F has a matching of size $\operatorname{vc}(F)/2$ and $\xi_{F,Q,G_{\bar{A}}}$ only depends on $G_{\bar{A}}$, we may condition on $G_{\bar{A}}$ and average over it to obtain that

$$\begin{aligned}
&\Pr_G\left[\left|\sum_{t_A\in\mathcal{T}_A} \chi_{F(t_A)}(G)\, \xi_{F,Q,G_{\bar{A}}}(t_A)\right| > s \,\middle|\, \boldsymbol{X}_r(F,Q,G)\right]\\
&\quad= \sum_{\substack{G_{\bar{A}}:\\ \boldsymbol{X}_r(F,Q,G)}} \Pr_{G\setminus G[V_A]}\left[G_{\bar{A}} \mid \boldsymbol{X}_r(F,Q,G)\right] \cdot \Pr_{G[V_A]}\left[\left|\sum_{t_A\in\mathcal{T}_A} \chi_{F(t_A)}(G)\, \xi_{F,Q,G_{\bar{A}}}(t_A)\right| > s\right] \ . \qquad (41)
\end{aligned}$$

For $r := 3n^{k-|A|}$, $m := n^{2-\delta/2}/4$, and $s := 6n^k p^{-|E(F)|}(m/n^2)^{\operatorname{vc}(F)/4}$ we find that

$$\Pr_{G[V_A]}\left[\left|\sum_{t\in Q}\sum_{H\in\mathcal{H}_F} \chi_{H(t)}(G)\right| > 6n^k\cdot p^{-|E(F)|}\cdot\left(\frac{m}{n^2}\right)^{\operatorname{vc}(F)/4} \,\middle|\, \boldsymbol{X}_r(F,Q,G)\right] < 2^{-m} \ . \qquad (42)$$

Since by Lemma 5.4 there are at most $\sum_{i=1}^{D/4} 2^{3i(i+\log k)} \le 2^{D(D+\log k)}$ many cores with vertex cover at most $D/4$ and by assumption it holds that $|\mathcal{Q}| \le 2^{n^{2-\delta}}$, it follows that the sum under consideration is bounded by

$$2^{D(D+\log k)} 2^{n^{2-\delta}} 2^{-m} = 2^{D(D+\log k) - n^{2-\delta/2}(1/4 - n^{-\delta/2})} = o(1) \ . \qquad \square$$

6 Concluding Remarks

We prove lower bounds on tree-like semantic proof systems operating over a limited number of distinct proof lines. These lower bounds are essentially optimal in a sense that the number of proof lines cannot be increased by much as otherwise there are semantic proof systems that can refute the entire family of formulas under consideration. As corollaries we obtain the first superpolynomial refutation length lower bounds on tree-like Frege refutation systems with bounded line size and the tree-like threshold proof system of polynomial degree. These lower bounds are established for different encodings of the clique formula. For the standard unary encoding, our main theorem yields essentially optimal average-case $n^{\Omega(D)}$ size lower bounds on tree-like cutting planes and tree-like resolution over parities refutations of the k-clique formula over Erdős-Rényi random graphs with maximum cliques of size D.

Given that we obtain lower bounds for proof systems that can refute all simple benchmark formulas we thereby establish that the proof method of constructing a pseudo-measure is independent of the hardness of these simple combinatorial principles. This shows that this proof paradigm may be able to yield lower bounds for even stronger proof systems. As a first step in this direction it would be interesting to see whether the lower bound methodology of constructing a pseudo-measure can be adapted to only hold for a certain class of *syntactic* derivation rules instead of the very generic semantic derivation rule considered in this paper. There are many further loose ends—let us mention two other directions we consider most interesting.

First off, is it possible to prove refutation length lower bounds on dag-like semantic $\mathcal{F}$ proof systems? This might be an avenue towards dag-like *Lovász–Schrijver* refutation length lower bounds. Could it even be that for every dag-like semantic $\mathcal{F}$ proof system there is a tree-like semantic $\mathcal{F}'$ proof system that efficiently simulates the former with $|\mathcal{F}'| = \mathrm{poly}(|\mathcal{F}|)$?

To prove cutting planes lower bounds for random $O(1)$-CNF formulas, we need to consider proof methods that separate deterministic from randomized communication models; the falsified clause search problem for random $O(1)$-CNF formulas has low-depth *randomized* communication protocols. As already pointed out in [dRPR23], the pseudo-measure μ can be used to prove a lower bound on the *deterministic* k-player number-in-hand communication complexity: consider the problem of finding a missing edge in the induced subgraph by a tuple $(v_1, \ldots, v_k)$, where player $i \in [k]$ is provided vertex v_i. The pseudo-measure μ shows that any deterministic communication protocol for this problem is of depth $\Omega(D \log n)$. As this problem can be solved in constant cost in the *randomized* model for small values of k, this provides a modest separation of these two models. Can the method of constructing a pseudo-measure be used to prove random $O(1)$-CNF cutting planes refutation size lower bounds?

Acknowledgements

We are grateful to Mika Göös, Dmitry Itsykson, Duri Andrea Janett, and Artur Riazanov for insightful discussions.

S.R. and D.E. received funding from the Knut and Alice Wallenberg Foundation grant KAW 2023.0116, ELLIIT, and the Swedish Research Council grant 2021-05104. Y.G. received funding from the Independent Research Fund Denmark grant 9040-00389B. K.R. is supported by the Swiss National Science Foundation Postdoc.Mobility fellowship P500-2_235298; part of this work was done while affiliated with EPFL. We also gratefully acknowledge that we have

benefited greatly from being part of the Basic Algorithms Research Centre (BARC) environment financed by the Villum Investigator grant 54451.

References

[AB09] Martin Anthony and Peter L Bartlett. *Neural network learning: Theoretical foundations*. cambridge university press, 2009. doi:10.1017/CBO9780511624216. 14

[AO18] Albert Atserias and Joanna Ochremiak. Proof complexity meets algebra. 20(1), December 2018. doi:10.1145/3265985. 1

[BFPJH21] Eric Blais, Renato Ferreira Pinto Jr., and Nathaniel Harms. VC dimension and distribution-free sample-based testing. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, page 504–517. Association for Computing Machinery, 2021. doi:10.1145/3406325.3451104. 14

[BIS07] Paul Beame, Russell Impagliazzo, and Ashish Sabharwal. The resolution complexity of independent sets and vertex covers in random graphs. *Computational Complexity*, 16(3):245–297, October 2007. doi:10.1007/s00037-007-0230-0. 8

[BPS07] Paul Beame, Toniann Pitassi, and Nathan Segerlind. Lower bounds for Lovász–Schrijver systems and beyond follow from multiparty communication complexity. *SIAM Journal on Computing*, 37(3):845–869, 2007. doi:10.1137/060654645. 3, 5, 7

[CEYZ20] Justin Chen, Christopher Eur, Greg Yang, and Mengyuan Zhang. Free resolutions of function classes via order complexes. *Advances in Applied Mathematics*, 120:102074, 2020. doi:10.1016/j.aam.2020.102074. 13

[CR79] Stephen A. Cook and Robert A. Reckhow. The relative efficiency of propositional proof systems. *Journal of Symbolic Logic*, 44(1):36–50, March 1979. Preliminary version in *STOC '74*. doi:10.2307/2273702. 1

[dRPR23] Susanna F. de Rezende, Aaron Potechin, and Kilian Risse. Clique is hard on average for unary Sherali–Adams. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 12–25, 2023. doi:10.1109/FOCS57990.2023.00008. 4, 15, 16, 17, 18, 20, 23

[dRPR24] Susanna F. de Rezende, Aaron Potechin, and Kilian Risse. Clique is hard on average for Sherali–Adams with bounded coefficients. *CoRR*, abs/2404.16722, 2024. doi:10.48550/ARXIV.2404.16722. 17, 18, 20, 21

[GHJ+24] Mika Göös, Alexandros Hollender, Siddhartha Jain, Gilbert Maystre, William Pires, Robert Robere, and Ran Tao. Separations in proof complexity and TFNP. *J. ACM*, 71(4), August 2024. doi:10.1145/3663758. 1

[GP18] Mika Göös and Toniann Pitassi. Communication lower bounds via critical block sensitivity. *SIAM Journal on Computing*, 47(5):1778–1806, 2018. doi:10.1137/16M1082007. 3, 5

[HKT24] Pavel Hubáček, Erfan Khaniki, and Neil Thapen. TFNP Intersections Through the Lens of Feasible Disjunction. In Venkatesan Guruswami, editor, *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, volume 287 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 63:1–63:24, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2024.63. 4

[IR21] Dmitry Itsykson and Artur Riazanov. Proof Complexity of Natural Formulas via Communication Arguments. In Valentine Kabanets, editor, *36th Computational Complexity Conference (CCC 2021)*, volume 200 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 3:1–3:34, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CCC.2021.3. 3, 5

[IS20] Dmitry Itsykson and Dmitry Sokolov. Resolution over linear equations modulo two. *Annals of Pure and Applied Logic*, 171(1):102722, 2020. doi:10.1016/j.apal.2019.102722. 7

[KM24a] Pravesh K. Kothari and Peter Manohar. An exponential lower bound for linear 3-query locally correctable codes. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*, pages 776–787. ACM, 2024. doi:10.1145/3618260.3649640. 1

[KM24b] Pravesh K. Kothari and Peter Manohar. Exponential lower bounds for smooth 3-LCCs and sharp bounds for designs. In *65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024*, pages 1802–1845. IEEE, 2024. doi:10.1109/FOCS61266.2024.00110. 1

[PZ24] Aaron Potechin and Aaron Zhang. Bounds on the Total Coefficient Size of Nullstellensatz Proofs of the Pigeonhole Principle. In Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson, editors, *51st International Colloquium on Automata, Languages, and Programming (ICALP 2024)*, volume 297 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 117:1–117:20, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ICALP.2024.117. 4

[Rec75] Robert A. Reckhow. *On the Lengths of Proofs in the Propositional Calculus*. PhD thesis, University of Toronto, 1975. doi:1807/100390. 1

[Sau72] N Sauer. On the density of families of sets. *Journal of Combinatorial Theory, Series A*, 13(1):145–147, 1972. doi:10.1016/0097-3165(72)90019-2. 6

[She72] Saharon Shelah. A combinatorial problem; stability and order for models and theories in infinitary languages. *Pacific Journal of Mathematics*, 41(1):247 – 261, 1972. doi:10.2140/pjm.1972.41.247. 6

[SSBD14] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014. doi:10.1017/CBO9781107298019. 12

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il