Hard Instances for Strong Semantic Proof Systems: CNFs for Depth-Four PC and Deep Instances for $\mathsf{AC}^0[2]$-Frege*

Yaroslav Alekseev
Technion

Svyatoslav Gryaznov
Imperial College London

Dmitry Itsykson
Ben-Gurion University

Toniann Pitassi
Columbia University

Robert Robere
McGill University

Iddo Tzameret
Imperial College London

Abstract

We establish a collection of new size lower bounds against fairly strong algebraic-Boolean propositional proof systems, from depth-3 semantic polynomial calculus for CNFs, up to semantic $\mathsf{AC}^0[2]$-Frege in the regime of deep formulas, as follows:

- $\Sigma\Pi\Sigma$-PC: Exponential lower bound against semantic $\Sigma\Pi\Sigma$-PC over $\mathbb{F}_2$, namely, polynomial calculus operating with depth-3 circuits with sum top gate, for the bit pigeonhole principle. We further obtain $2^{\Omega(n)}$ lower bounds for random 3-CNFs on n variables and for the bit and functional pigeonhole principles with $n+1$ pigeons and n holes.
- $\Sigma\Pi\Sigma\Pi$-PC: Superpolynomial lower bounds for semantic $\Sigma\Pi\Sigma\Pi$-PC over $\mathbb{F}_2$, with unrestricted bottom fan-in. The hard instances include the weak unary pigeonhole principle and lifting of random CNFs of logarithmic width. For the weak pigeonhole principle with $2n$ pigeons and n holes, the lower bound is $\exp(\Omega((\log n)^2/\log\log n))$. This further settles the open problem of establishing superpolynomial lower bounds against $\mathsf{Res}(\log)$ for this principle. When the $\Sigma\Pi\Sigma\Pi$-PC bottom fan-in is logarithmic, we obtain exponential lower bounds for random CNFs of logarithmic width.
- Depth-d $\mathsf{AC}^0[2]$-Frege: For every fixed $d \geq 2$, we construct unsatisfiable families of formulas of depth d that require superpolynomial-size semantic depth-d $\mathsf{AC}^0[2]$-Frege refutations, but have polynomial-size semantic depth-$(d+1)$ $\mathsf{AC}^0[2]$-Frege refutations. This provides the first unconditional hard instances against arbitrary depth $\mathsf{AC}^0[2]$-Frege.

The arguments build and expand on the recent AI-generated proof of an exponential lower bound against resolution over parities by Braun [16]. We distil Braun's implicit *annihilating polynomial* argument into an abstract framework based on standard pseudo-reduction operators, and augment it with new techniques. We show how rank lower bounds for these operators imply proof-size lower bounds. The polynomial approximation method of Razborov and Smolensky [60, 64] extends the framework to bounded bottom degree. Lifting and random restrictions remove the restriction on bottom fan-in. For the $\mathsf{AC}^0[2]$-Frege result, we combine polynomial approximation with the common annihilating polynomial and a lifting construction based on constant-depth formulas for the coin problem.

AI Methodology. The authors used OpenAI models, including GPT-5.6 to GPT-6.1, as interactive tools to simplify existing proofs, explore and develop mathematical arguments. The paper was written and verified by the authors. The authors take full responsibility for the content and correctness of the paper.

*S.G., T.P., R.R., and I.T. would like to thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the programme Logical Foundations of Complexity Theory, where work on this paper was undertaken. This work was supported by EPSRC grant EP/Z000580/1. Yaroslav Alekseev was supported by Israel Science Foundation grant 507/24. Dmitry Itsykson was supported by European Research Council grant 949707. Robert Robere was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC). The work of Svyatoslav Gryaznov and Iddo Tzameret received funding from the European Research Council under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 101002742, EPRICOT project). Iddo Tzameret was also supported by the Engineering and Physical Sciences Research Council under grant EP/Z534158/1.

Contents

1 Introduction **1**
- 1.1 Our Results . . . 1
  - 1.1.1 Semantic Algebraic Proof Systems . . . 2
  - 1.1.2 Statement of Results . . . 4
- 1.2 Overview of the Lower Bound Arguments . . . 5
  - 1.2.1 Depth-Three Lower Bounds by Annihilating Polynomials . . . 6
  - 1.2.2 Bounded-Bottom-Degree Lower Bounds by the Approximation Method . . . 7
  - 1.2.3 Depth-Four Lower Bounds by Lifting . . . 8
  - 1.2.4 Depth Hierarchy . . . 8
- 1.3 Concurrent results . . . 10
- 1.4 Paper Outline . . . 13

2 Preliminaries **13**
- 2.1 Propositional proof systems . . . 14
- 2.2 Semantic proof systems . . . 16
- 2.3 Simulation of $\mathsf{Res}(\oplus)$ . . . 18
- 2.4 Simulation of $\mathsf{Res}(\mathsf{PC})$ . . . 19

3 Warm-Up: $\Sigma\Pi\Sigma$-PC Lower Bound for the Bit Pigeonhole Principle **20**

4 Abstract Lower Bound Framework: Degree Reduction **24**
- 4.1 Proof of the $\Sigma\Pi\Sigma$-PC size bound . . . 27
- 4.2 Proof of the $\Sigma\Pi\Sigma\Pi^{[w]}$-PC size bound . . . 29
  - 4.2.1 Proof of the approximation lemma. . . . 31
- 4.3 Proof of the $\mathsf{AC}^0[2]$-Frege size bound . . . 34

5 Rank Bounds for Existing Pseudo-Reduction Operators **37**
- 5.1 The expansion pseudo-reduction operator . . . 37
- 5.2 Rank estimates . . . 39

6 Lower Bounds for Semantic $\Sigma\Pi\Sigma$-PC **40**
- 6.1 Random 3-CNFs . . . 41
- 6.2 Functional pigeonhole principle . . . 43
- 6.3 Semantic $\Sigma\Pi\Sigma\Pi^{[2]}$-PC is strictly stronger than $\Sigma\Pi\Sigma$-PC . . . 48

7 Lower Bounds for Semantic $\Sigma\Pi\Sigma\Pi$-PC **50**
- 7.1 Random $O(\log n)$-CNFs in semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC . . . 50
- 7.2 Graph PHP in semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC . . . 52
- 7.3 Lifting to semantic $\Sigma\Pi\Sigma\Pi$-PC . . . 54
  - 7.3.1 Representations and the size measure . . . 54
  - 7.3.2 The selector formula and its projections . . . 55
  - 7.3.3 Compression and proof lifting . . . 56
  - 7.3.4 Application to random $O(\log n)$-CNFs . . . 58
- 7.4 Sparsification: PHP in semantic $\Sigma\Pi\Sigma\Pi$-PC . . . 58

8 Lower Bounds for Semantic $\mathsf{AC}^0[2]$-Frege **61**
- 8.1 The lifting operation . . . 62
  - 8.1.1 Why the gadget needs the extra depth . . . 63
- 8.2 Degree bounds on the lifted formula . . . 63
- 8.3 Gadget construction . . . 66
- 8.4 The size lower bound for lifted formulas . . . 69

A Proof of the Pigeon-Degree Lower Bound **70**

B Proof of the Right-Degree Reduction Lemma **71**

C Unsatisfiability and Expansion of Random Logarithmic-Width CNFs **72**

D Boundary Expanders for the Pigeonhole Principle **73**

References **74**

1 Introduction

Proving superpolynomial lower bounds for $\mathsf{AC}^0[2]$-Frege is considered one of the central open problems in the technical frontiers of propositional proof complexity. This proof system has lines that are bounded-depth Boolean formulas with parity gates. While strong lower bounds for the corresponding circuit class have been known since the work of Razborov and Smolensky [60, 64], extending these lower bound techniques to proofs has remained a significant challenge since the 1990s (cf. [11, 18]). A natural approach to this problem is to study fragments of $\mathsf{AC}^0[2]$-Frege which retain some of its algebraic structure, but impose further restrictions on the formulas that can occur in a proof. Two prominent examples are *polynomial calculus over* $\mathbb{F}_2$ (PC_2), whose lines are polynomials represented as sums of monomials, and *resolution over parities* ($\mathsf{Res}(\oplus)$), whose lines are disjunctions of linear equations over $\mathbb{F}_2$.

In this paper, we prove size lower bounds for a collection of algebraic proof systems which generalise *both* of these proof systems, as well as for depth-d $\mathsf{AC}^0[2]$-Frege (in the regime where the hard formulas are of depth-d). For our algebraic proof systems, their lines are represented by bounded-depth algebraic circuits over $\mathbb{F}_2$, and the inference rules are *semantic* (which are stronger than standard syntactic inference): any polynomial equation which follows from two earlier equations over the Boolean cube may be derived. We collectively refer to these systems as *semantic bounded-depth polynomial calculus*, and if $\mathcal{C}$ is a class of algebraic circuits, we write $\mathcal{C}$-PC for short.

We obtain exponential lower bounds for semantic depth-three polynomial calculus, including for random 3-CNFs and the bit and functional pigeonhole principles, and superpolynomial lower bounds for semantic depth-four polynomial calculus with unrestricted bottom fan-in. The latter bounds apply to the weak unary pigeonhole principle, and to a family of CNFs obtained by lifting random formulas.

The starting point for our lower bounds is a recent breakthrough of Braun [16] who proved exponential lower bounds for $\mathsf{Res}(\oplus)$. This system is the modulo-two variant of resolution over linear equations introduced by Raz and Tzameret [57], and was extensively studied by Itsykson and Sokolov [44, 45]. Allowing disjunctions of linear equations in place of ordinary clauses gives a natural way of combining resolution with reasoning about parity constraints. However, despite substantial progress on tree-like refutations [44, 45, 32, 50, 36, 43, 12, 21, 37, 20, 63], regular refutations [15, 26, 5], and refutations of bounded depth [5, 27, 3, 41, 14, 28, 42, 62], proving a superpolynomial lower bound for unrestricted DAG-like $\mathsf{Res}(\oplus)$ had remained open. Braun's work resolved this problem by proving that every $\mathsf{Res}(\oplus)$ refutation of the bit pigeonhole principle with $n+1$ pigeons and $n = 2^\ell$ holes requires $2^{\Omega(n/\log^2 n)}$ lines.

1.1 Our Results

Summary. This work aims to provide both conceptual and methodological-technical contributions.

1. *Algebraic systems.* Instead of focusing on $\mathsf{Res}(\oplus)$ and its extensions, we show that the annihilating polynomial technique from [16] together with our new extensions fit naturally into constant-depth algebraic proof systems.
2. *Semantic inference and barriers to lower bounds.* Similarly, we show that these techniques apply already to the stronger notion of semantic proof systems instead of the usual syntactic ones. This strengthens our result, while also revealing a potential limitation on the applicability of the techniques as explained below.
3. *Exposition.* We give a simpler presentation of Braun's argument, while observing that it applies to the stronger system of semantic $\Sigma\Pi\Sigma$ polynomial calculus. The central idea implicit in Braun's proof is to use a low-degree *annihilating polynomial* which functions in analogy to a restriction argument: multiplying every line by this polynomial eliminates the high-degree product terms in a purported small proof (reducing size to degree lower bounds).
4. *Methodology.* We formulate an abstract version of the annihilating polynomial argument in terms of the *rank of a pseudo-reduction operator*. Pseudo-reductions are familiar objects from polynomial calculus degree lower bounds [23, 59, 2]. We combine this method with Braun's dimension argument, using large rank to obtain size lower bounds. Combining this formulation with the approximation

method of Razborov and Smolensky [60, 64], along with techniques from *lifting theorems*, allows us to move beyond $\mathsf{Res}(\oplus)$ to proofs operating with depth-four algebraic circuits. We choose a common annihilating polynomial so that, after multiplication, each approximation equals the corresponding original polynomial. This preserves every sound binary inference. We show that the same approximation and annihilation argument yields lower bounds for depth-d semantic $\mathsf{AC}^0[2]$-Frege on depth-d formulas. This is achieved by turning a putative short refutation of a lifted formula into a polynomial calculus refutation of the original (unlifted) CNF, obtaining a contradiction with its degree lower bound.

5. *CNF hardness up to the next-depth semantic collapse.* In the alternating sum/product hierarchy with a sum gate at the root, depth five already admits polynomial-size semantic refutations of every unsatisfiable CNF. Our unrestricted depth-four lower bounds reach the depth immediately below this collapse. They therefore concern the largest depth at which superpolynomial semantic CNF lower bounds can hold in this hierarchy.
6. $\mathsf{AC}^0[2]$*-Frege lower bounds at every fixed depth.* For every fixed $d \geq 2$, we construct unsatisfiable families of depth-d formulas that require superpolynomial-size semantic depth-d $\mathsf{AC}^0[2]$-Frege refutations. The proof lines may use the full basis of AND, OR, and parity gates, with unbounded fan-in. The same instances have polynomial-size semantic depth-$(d+1)$ refutations using only AND and OR gates. Thus we obtain a separation between consecutive semantic proof depths; the input formulas have depth d and depend on the chosen depth.

We first describe the proof systems that we consider, and then state precisely our main lower bounds. Throughout the paper, the ground field is $\mathbb{F}_2$, and all semantic inferences use at most two premises (this is important, as mentioned in the subsection below).

1.1.1 Semantic Algebraic Proof Systems

To discuss our systems, it is convenient to identify a polynomial with the Boolean function that it computes. Formally, this means working in the quotient ring

$$\mathcal{B}_n = \mathbb{F}_2[x_1, \ldots, x_n]/\langle x_1^2 - x_1, \ldots, x_n^2 - x_n \rangle.$$

Every element of this ring has a unique multilinear representative. A polynomial P occurring as a proof line will be interpreted as representing the equation $P = 0$, and a refutation ends with the constant polynomial 1. We encode a clause by the polynomial which is 1 precisely on the assignments falsifying it. For example, the clause $x \vee \neg y$ is represented by $(1 - x)y$.

Let $\mathcal{C}$ be a class of algebraic circuits. A *semantic* $\mathcal{C}$-PC *refutation* is a sequence of polynomials represented by circuits from $\mathcal{C}$, in which each polynomial is an axiom or follows from at most two earlier polynomials by a sound semantic inference. In other words, we may derive R from P and Q whenever

$$P(x) = Q(x) = 0 \implies R(x) = 0 \qquad \text{for every } x \in \{0,1\}^n.$$ [1]

The *size* of a refutation is the number of distinct rooted subcircuits appearing anywhere in its lines, allowing sharing across proof lines. This definition separates two aspects of an algebraic proof: the representation of its polynomials, and the rules used to derive them. We restrict the first while allowing every sound binary inference for the second.

For $\mathcal{C} = \Sigma\Pi$, we obtain semantic polynomial calculus with the usual sparse representation of polynomials. The main systems we consider are $\Sigma\Pi\Sigma$-PC and $\Sigma\Pi\Sigma\Pi$-PC, which allow one or two further layers. A line in $\Sigma\Pi\Sigma$-PC is a sum of products of affine forms. A line in $\Sigma\Pi\Sigma\Pi$-PC is a sum of products of polynomials, each of which is itself represented as a sum of monomials. We write $\Sigma\Pi\Sigma\Pi^{[w]}$-PC when the bottom product gates have fan-in at most w, or, equivalently, when each of these bottom polynomials is represented using monomials of degree at most w. In particular, $w = 1$ gives the depth-three system.

[1] For context, in *syntactic* PC we can only infer $x_i \cdot P(x) = 0$, for x_i a variable, from $P(x) = 0$ and $\alpha \cdot P(x) + \beta \cdot Q(x) = 0$ from $P(x) = Q(x) = 0$, for α, β scalars.

These systems simulate several familiar proof systems. To see the connection with $\mathsf{Res}(\oplus)$, consider a linear clause

$$(\ell_1 = b_1) \vee \cdots \vee (\ell_t = b_t).$$

Over $\mathbb{F}_2$, this clause is equivalent to the polynomial equation

$$\prod_{j=1}^{t} (\ell_j + b_j) = 0.$$

Hence each line of a $\mathsf{Res}(\oplus)$ proof becomes a single product of affine forms, and every sound inference remains sound after the translation. This gives a simulation of semantic $\mathsf{Res}(\oplus)$ by semantic $\Sigma\Pi\Sigma$-PC with at most one product term per original proof line; see Theorem 2.3. More generally, translating a disjunction of polynomial equations into the product of its polynomials gives a simulation of semantic $\mathsf{Res}(\mathsf{PC}_w)$ by semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC (because the bottom two layers $\Sigma\Pi^{[w]}$ can express polynomials of degree w, while the disjunctions in the former system are expressed by the Π outer layer; see Theorem 2.4). The unrestricted depth-four system also polynomially simulates $\mathsf{Res}(\log)$; here the width of a term is logarithmic in the size of the proof (instead of the number of variables) (Proposition 7.15).

Allowing all sound binary inferences, instead of the standard syntactic ones, strengthens the proof system. A lower bound for the semantic system therefore holds for every sound syntactic system with the same representations and corresponding binary syntactic rules. Note that we impose no requirement that the validity of an inference be efficiently checkable, hence the semantic proof systems are not Cook–Reckhow systems [25]. In particular, it is not possible to verify an application of an inference rule in polynomial time. This additional power has the following consequence: semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC efficiently refutes *every* unsatisfiable system of degree-w polynomials, provided that the input polynomials have small sparse representations; and every unsatisfiable CNF has a polynomial-size semantic $\Sigma\Pi\Sigma\Pi\Sigma$-PC refutation, regardless of its clause width.[2]

Specifically, for the lower bounds we establish this means the following:

- Our $\Sigma\Pi\Sigma\Pi$-PC lower bounds hold one layer below a system in which every unsatisfiable CNF has a short refutation.

Relation to semantic $\mathsf{AC}^0[2]$*-Frege.* The $\mathsf{AC}^0[2]$-Frege results employ the close connection between the algebraic systems and $\mathsf{AC}^0[2]$-Frege. Over $\mathbb{F}_2$, addition computes parity, multiplication computes conjunction, and negation is given by $1 + x$. Algebraic representations of this kind therefore provide a natural way of studying bounded-depth reasoning with parity. Our semantic $\Sigma\Pi\Sigma$-PC and $\Sigma\Pi\Sigma\Pi$-PC are simulated line by line by semantic depth-three and depth-four $\mathsf{AC}^0[2]$-Frege, respectively, by reading sums as parities, products as conjunctions, and complementing the output.

Similarly to semantic $\Sigma\Pi\Sigma\Pi\Sigma$-PC above, every unsatisfiable family $\mathcal{F}$ of axioms of depth strictly less than d has a polynomial-size semantic depth-d $\mathsf{AC}^0[2]$-Frege refutation. Hence:

- For semantic $\mathsf{AC}^0[2]$-Frege the only way to prove a lower bound is to consider hard formulas of *maximum* allowed depth, as we do.

Related (dynamic, i.e., line-by-line) *syntactic* systems operating with algebraic circuits were studied in numerous works [34, 58, 57, 65, 19, 39]. Related static (i.e., one-shot) systems include Pitassi's algebraic proof system [56] and the *Ideal Proof System* (IPS) of Grochow and Pitassi [35]. Specifically,

[2]To see why, suppose we have a family of polynomials $\mathcal{F} = \{f_1, \ldots, f_t\}$ over $\mathbb{F}_2$ without a common zero (for example, we can take a family obtained by translating an unsatisfiable CNF formula). For $A \subseteq [t]$ define

$$Q_A = 1 + \prod_{i \in A} (1 + f_i),$$

so that the equation $Q_A = 0$ asserts that every f_i with $i \in A$ is zero. Hence, for any $A, B \subseteq [t]$, $Q_{A \cup B}$ follows semantically from Q_A and Q_B. Building these conjunctions along a balanced binary tree gives a refutation with at most $2t - 1$ lines, since $Q_{[t]}$ is the constant function 1 when the input is unsatisfiable. Its total circuit size is polynomial in the number of monomials occurring in $\mathcal{F}$. In particular, unsatisfiable CNFs of width at most w have short semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutations when w is constant.

Grochow and Pitassi [35, Theorem 3.5] show that depth-d $\mathsf{AC}^0[2]$-Frege refutations with the usual syntactic inference rules can be transformed into IPS refutations over $\mathbb{F}_2$ of polynomially related size and depth $O(d)$. Hakoniemi, Limaye and Tzameret [38] established constant-depth IPS lower bounds; however the lower bounds hold only when the individual degree of the proof is $O(\log \log n)$, for which $\mathsf{AC}^0[2]$-Frege simulation breaks.

For every fixed depth, Buss, Kołodziejczyk, and Zdanowski [19, Theorem 43] give a quasipolynomial simulation of the usual $\mathsf{AC}^0[2]$-Frege refutations of CNFs by a syntactic system whose lines are disjunctions of conjunctions of sparse polynomial equations. Encoding each whole line by a single polynomial equation gives a depth-five $\Pi\Sigma\Pi\Sigma\Pi$ representation over $\mathbb{F}_2$, one algebraic layer deeper than our $\Sigma\Pi\Sigma\Pi$-PC lines. The syntactic rules matter here: allowing arbitrary semantic inference in their target system gives short refutations of every unsatisfiable CNF, as above.

1.1.2 Statement of Results

Our lower bounds are summarised in Table 1. For any proof π in $\mathcal{C}$-PC, where $\mathcal{C} \in \{\Sigma\Pi\Sigma, \Sigma\Pi\Sigma\Pi^{[w]}\}$, we let $S(\pi)$ denote the number of *distinct non-zero top-product terms* occurring in the refutation, and we identify terms which compute the same Boolean function. This is a lower bound on total circuit size. For $\Sigma\Pi\Sigma\Pi$-PC, our bounds concern total circuit size, since the projection arguments must also account for the distinct monomials occurring near the bottom of the circuits.

No.	Instance	Proof system	Lower bound
1	Bit BPHP_n^{n+1}, $n = 2^\ell$	$\Sigma\Pi\Sigma$-PC	$2^{\Omega(n)}$
2	Functional FPHP_n^{n+1}	$\Sigma\Pi\Sigma$-PC	$2^{\Omega(n)}$
3	Random $\Theta(n)$-size 3-CNFs	$\Sigma\Pi\Sigma$-PC	$2^{\Omega_\alpha(n)}$
4	Random $\Theta(n^2)$-size $\Theta(\log n)$-CNF	$\Sigma\Pi\Sigma\Pi^{[w]}$-PC	$\exp\left(\Omega\left(\frac{n}{wC_1^w}\right)\right)$
5	Graph PHP, $2n$ pigeons, n holes	$\Sigma\Pi\Sigma\Pi^{[w]}$-PC	$\exp\left(\Omega\left(\frac{n}{w\Delta(C_2\Delta^2)^w}\right)\right)$
6	Lifted random $\Theta(\log n)$-CNF	$\Sigma\Pi\Sigma\Pi$-PC	$\exp\left(\Omega((\ln n)^2)\right)$
7	Unary PHP_n^{2n}	$\Sigma\Pi\Sigma\Pi$-PC	$\exp\left(\Omega(\frac{(\ln n)^2}{\ln\ln n})\right)$
8	Unary PHP_n^{2n}	$\mathsf{Res}(\log)$	$\exp\left(\Omega(\frac{(\ln n)^2}{\ln\ln n})\right)$
9	Lifted depth-d formulas	Depth-d $\mathsf{AC}^0[2]$-Frege	$\exp(\Omega_d((\log N)^{d/(d-1)}))$

Table 1: Lower bounds. The algebraic and Frege systems are semantic; the algebraic systems are over $\mathbb{F}_2$. Rows 1–5 bound $S(\pi)$; rows 6, 7, and 9 bound total circuit size. Row 8 bounds $\mathsf{Res}(\log)$ proof size and follows from row 7 (Corollary 7.16).[3] Random-instance bounds hold with high probability. In row 3, α is a sufficiently large fixed clause density. In row 4, $1 \le w \le c \log_2 n$. In row 5, $\Delta = \Theta(\ln n)$ and $1 \le w \le \ln n / \ln\ln n$. The constants $C_1, C_2 > 1$ and $c > 0$ are absolute. The parameter n counts holes in the pigeonhole instances and variables in the unlifted random formulas. In row 9, $d \ge 2$ is fixed, N is the total circuit size of the input family.

We give a few remarks on these results. Our first contribution is a short proof of Braun's lower bound for the bit pigeonhole principle. More precisely, for $n = 2^\ell$ and $m > n$, every semantic $\Sigma\Pi\Sigma$-PC refutation of BPHP_n^m requires

$$2^{\Omega(n^2/(m\ell^2))}$$

distinct non-zero product terms (Theorem 3.2). For $n < m \le 2n$, this recovers Braun's exponent in a stronger system. Later, by applying the general framework to the functional pigeonhole principle and transferring the resulting bound by an affine substitution, we improve the bound for BPHP_n^{n+1} to $2^{\Omega(n)}$ (Corollaries 6.13 and 6.14).

Next, we apply our improved method to random formulas. For sufficiently large constant α, a random 3-CNF with n variables and αn independently and uniformly sampled clauses is unsatisfiable with high probability, and we show it requires $2^{\Omega_\alpha(n)}$ product terms in semantic $\Sigma\Pi\Sigma$-PC (Theorem 6.5).

[3] Ondřej Ježil and Dimitrios Tsintsilidas obtained this result independently of us [46].

To obtain depth-four lower bounds, we must increase the clause width. If $\ell = \lceil \log_2 n \rceil$ and F_n consists of $n2^\ell$ independently and uniformly sampled width-ℓ clauses, then with high probability it requires

$$\exp\left(\Omega\left(\frac{n}{wC_1^w}\right)\right)$$

top product terms in semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC, simultaneously (i.e., for the same F_n) for every $1 \le w \le c \log_2 n$ (Theorem 7.2). For constant w this bound is exponential in n. The same family therefore remains hard as we increase the bottom fan-in through a logarithmic range.

Third, we obtain lower bounds *without any restriction* on bottom fan-in by lifting the formula with a "selector" gadget. For random F_n, applying a lift to F_n gives an unsatisfiable CNF with $2n^2$ variables and polynomially many literal occurrences which requires total circuit size $\exp\left(\Omega((\ln n)^2)\right)$ in semantic $\Sigma\Pi\Sigma\Pi$-PC (Theorem 7.12). For the unary pigeonhole principle, a more direct approach, namely, random restriction, suffices: we sparsify the complete pigeon–hole graph, and use the graph lower bound in the fifth row of the table. This yields the lower bound

$$\exp\left(\Omega\left(\frac{(\ln n)^2}{\ln \ln n}\right)\right)$$

for semantic $\Sigma\Pi\Sigma\Pi$-PC refutations of PHP_n^{2n} (Theorem 7.14). Since Res(log) is polynomially simulated by this system, we also obtain a superpolynomial lower bound for Res(log) refutations of the weak pigeonhole principle. The required proof size is at least $\exp\left(\Omega\left(\frac{(\ln n)^2}{\ln \ln n}\right)\right)$ (Corollary 7.16).

Finally, the functional and non-functional pigeonhole principles behave differently in our semantic systems. Although FPHP_n^{n+1} requires exponential size at depth three, it has semantic $\Sigma\Pi\Sigma\Pi^{[2]}$-PC refutations of size $O(n^3 \log n)$ (Theorem 6.18). The functionality axioms allow us to replace the wide pigeon clauses by linear equations asserting that each row has sum one. The resulting unsatisfiable system has degree at most two, and hence can be refuted by the conjunction construction described above. Together with the lower bound, this gives an exponential separation between semantic depth-three polynomial calculus and semantic depth-four polynomial calculus with bottom fan-in two. It also explains why we do not add functionality axioms to the weak pigeonhole principle in our depth-four lower bound.

The first eight rows of Table 1 concern CNF formulas, represented by their standard clause polynomials. The final row gives lower bounds for the full semantic $\mathsf{AC}^0[2]$-Frege system at every fixed depth. For every fixed $d \ge 2$, there are unsatisfiable families of formulas of depth at most d and total circuit size N whose semantic depth-d refutations require size

$$\exp\left(\Omega_d\left((\log N)^{d/(d-1)}\right)\right).$$

The same families have polynomial-size semantic depth-$(d+1)$ refutations using only AND and OR gates (Theorem 8.1). This gives a separation between consecutive semantic proof depths even when the smaller depth permits parity gates. The depth of the instance is essential: every unsatisfiable CNF has a short semantic refutation of Boolean depth two, obtained by a successive conjunction of their clauses and then using semantic inference.

1.2 Overview of the Lower Bound Arguments

The common idea behind our lower bounds is a reduction from the size of an algebraic proof to the degree of a polynomial calculus derivation. The familiar version of this strategy applies a random restriction to a small proof, hoping to eliminate its complicated lines while leaving a formula which still requires large degree. Here we replace the restriction with *multiplication by a carefully chosen low-degree polynomial*. This allows us to simplify a proof without assigning values to its variables.

There is an important difference between the two operations. A restriction preserves the final contradiction 1, whereas multiplication by a polynomial f changes the final line to f. Hence we must ensure not only that multiplication simplifies the proof, but also that f itself cannot be derived in low degree. Pseudo-reduction operators provide a convenient way of imposing both requirements at once. The additional property we need, beyond the usual degree lower bound, is that the operator has large rank on a suitable space of low-degree polynomials.

1.2.1 ΣΠΣ Lower Bounds by Annihilating Polynomials

We first explain the argument for ΣΠΣ-PC, whose lines are *sums of products of affine forms*. This part of the method is already implicit in Braun's proof [16, Section 6.1]. Suppose we are given a semantic ΣΠΣ-PC refutation with only a small number of distinct product terms, and consider one such term

$$T = \prod_j \ell_j,$$

where each ℓ_j is affine over $\mathbb{F}_2$. The polynomial T is the indicator of the set

$$A_T = \{x : \ell_j(x) = 1 \text{ for every } j\}.$$

If T is non-zero, then A_T is an affine subspace, and moreover, the degree of T in the Boolean quotient is exactly the codimension of this subspace.

Our goal is to find a low-degree polynomial f which *annihilates* (i.e. sets to 0) every term in the proof of sufficiently high degree. Since T is zero outside A_T, we have

$$fT = 0 \mod \langle x_i^2 - x_i \rangle_i \quad \iff \quad f|_{A_T} = 0.$$

Notice that this equality is in the Boolean quotient, rather than the formal polynomial ring.

The useful feature of a high-degree product term is that its support has high codimension. So, suppose that $\operatorname{codim}(A_T) \geq r$, and let U be a vector space of candidate polynomials of degree at most k. We can parametrise A_T by at most $n - r$ free coordinates, expressing every other coordinate as an affine function of them. Substituting these expressions into $f \in U$, and then multilinearising, does not increase its degree. Thus $f|_{A_T}$ belongs to a vector space of dimension at most

$$H_{n-r}(k) := \sum_{j=0}^{\min\{k,n-r\}} \binom{n-r}{j}.$$

Multiplying every line of the proof by f now removes (or *annihilates*) all terms of degree at least r. Every remaining term has degree less than r, so every transformed line has degree at most $k + r - 1$. The semantic inferences remain sound. Indeed, on an assignment where $f = 0$ all three transformed lines vanish, while on an assignment where $f = 1$ the transformed inference is the original one. The transformed proof is therefore a low-degree semantic *derivation* of f from the multiplied axioms (because the terminal line of the original refutation was 1).

It remains to choose U so that this is impossible, namely every $f \in U$ requires high degree derivation. For this purpose, let $\mathcal{R}$ be a *pseudo-reduction operator* defined on polynomials of degree at most B. This is a linear map which sends the input axioms to zero, and whose kernel is closed under multiplication by a variable whenever the resulting degree stays at most B. These properties imply, by induction on a polynomial calculus derivation, that every polynomial derivable in degree at most B lies in $\ker(\mathcal{R})$. Such operators are a standard tool in polynomial calculus degree lower bounds [23, 59, 2, 54]. Here we need more than a non-zero operator: its rank on low-degree polynomials must be large. This lets us combine the operator method with Braun's dimension argument to find an annihilating polynomial f with $\mathcal{R}(f) \neq 0$.

Choose U to be a complement to the kernel of $\mathcal{R}$ on the polynomials of degree at most k. Then every non-zero $f \in U$ satisfies $\mathcal{R}(f) \neq 0$, and

$$\dim U = \operatorname{rank}(\mathcal{R}_{\leq k}).$$

Hence the same dimension count finds an annihilator which is also hard to derive, provided that this rank is sufficiently large. If the input axioms have degree at most d, the multiplied axioms can be derived in degree at most $k+d$. Finally, every semantic inference over $\mathbb{F}_2$ between degree-D polynomials has an ordinary polynomial calculus simulation of degree at most $3D$ [19]. Taking $B = 3(k+\max\{d, r\})$ therefore gives a contradiction: the transformed proof yields a degree-B derivation of f, but $\mathcal{R}(f) \neq 0$.

This proves a size lower bound of the form

$$S(\pi) \geq \frac{\operatorname{rank}(\mathcal{R}_{\leq k})}{H_{n-r}(k)};$$

we refer to Theorem 4.4 for the precise parameter conditions. In other words: to rule out a *low-degree* refutation, it is enough that the pseudo-reduction operator is non-zero. But, to rule out a small depth-three proof, we need to show more: namely, that there are enough independent polynomials to survive under the operator that some non-zero combination also satisfies all the annihilation conditions. Equivalently, considered as a linear operator, we need that the pseudo-reduction has *large rank*.

Luckily, we are able to prove these rank estimates, and the particular estimates our applications use come from different sources. For the bit pigeonhole (warm-up, Section 3), we use the pigeon-degree lower bound of Razborov and Impagliazzo, Pudlák, and Sgall [59, 40]. For random formulas and graph pigeonhole principles, we use the expansion method of Mikša and Nordström [54], which generalises the method of Alekhnovich and Razborov [2]. In these applications, expansion supplies the pseudo-reduction operator, and counting suitable monomials which it fixes supplies its rank. Thus the new task is to extract sufficiently many independent surviving polynomials from the degree lower bound argument.

1.2.2 $\Sigma\Pi\Sigma\Pi^{[w]}$ Lower Bounds by the Approximation Method

For depth-four circuits, a top product term has the form

$$T = \prod_{j \in J} g_j,$$

where each g_j has degree at most w. Its support is now the common solution set of polynomial equations of degree at most w, rather than an affine subspace. We can no longer parametrise this set by affine functions of a smaller number of coordinates, so the dimension estimate from the previous argument does not directly apply.

We address this difficulty by combining the annihilating polynomial argument with the approximation method of Razborov and Smolensky [60, 64]. The first observation is that a long product can be approximated by a product of a small number of linear combinations of its factors. More precisely, choose

$$q_1, \ldots, q_r \in \operatorname{span}_{\mathbb{F}_2}\{1 + g_j : j \in J\},$$

and define

$$a = \prod_{i=1}^{r} (1 + q_i).$$

The degree of a is at most rw. Whenever $T(x) = 1$, every $g_j(x)$ is one, so all the $q_i(x)$ vanish and $a(x) = 1$. Thus $aT = T$: any error occurs at an input where $T = 0$ and $a = 1$, and so we can measure the errors by the function $T + a$.

The usual approximation argument chooses the q_i at random and bounds the probability of such an error at any fixed input. For our purpose, we need a different kind of control. We want one polynomial f which vanishes on the errors for *all* product terms in the proof, while still satisfying $\mathcal{R}(f) \neq 0$. To generalise the argument from the previous section, the relevant quantity is then the number of independent linear conditions imposed on a low-degree f by the equation $f(T + a) = 0$. Since the product $f(T + a)$ depends linearly on the coefficients of f, this number is the dimension of the space

$$(T + a)\mathcal{B}_n^{\leq k} = \{f(T + a) : f \in \mathcal{B}_n^{\leq k}\}.$$

We show that the q_i can be chosen so that this dimension decreases exponentially with r. In the Boolean cube, writing $H_n(k) = \dim \mathcal{B}_n^{\leq k}$, the bound takes the form

$$\dim\Big((T + a)\mathcal{B}_n^{\leq k}\Big) \leq H_n(k)(1 - \alpha)^r,$$

where

$$\alpha = \frac{H_n(k-w)}{2H_n(k)}.$$

The same statement holds in the structured quotients needed for the pigeonhole applications, with the corresponding dimension function $H_{\mathcal{F}_0}$; see Lemma 4.9.

We can now repeat the dimension argument from the previous section, using the same space U complementary to $\ker(\mathcal{R}_{\leq k})$. Choose an approximation a_T for each product term T, and suppose that each space $(T + a_T)\mathcal{B}_n^{\leq k}$ has dimension at most λ. The equations

$$f(T + a_T) = 0$$

for all product terms impose at most $S(\pi)\lambda$ homogeneous linear conditions on $f \in U$. If this is smaller than $\dim U = \operatorname{rank}(\mathcal{R}_{\leq k})$, there is a non-zero solution $f \in U$, and hence $\mathcal{R}(f) \neq 0$. Consequently, $fT = fa_T$ has degree at most $k + rw$ for every product term. Multiplication by f again converts the original refutation into a low-degree derivation of f, contradicting the pseudo-reduction operator. This gives the general bounds in Theorem 4.5 and Corollary 4.6.

In this formulation, approximation and annihilation play complementary roles. The approximation produces a low-degree polynomial whose errors impose relatively few linear conditions, while the annihilator makes all these approximations exact throughout the proof. This is what allows the argument to accommodate arbitrary sound binary inferences, without having to preserve a particular syntactic derivation rule.

1.2.3 $\Sigma\Pi\Sigma\Pi$ Lower Bounds by Lifting

The preceding argument gives useful bounds when the bottom degree w is small. To obtain lower bounds for unrestricted depth-four circuits, we use projections which simplify the bottom polynomials while preserving a hard instance. We perform two distinct versions of this step: first, a general selector lift that applies to random formulas, and a second direct sparsification argument for the weak pigeonhole principle.

In the selector construction, we replace each variable of a base CNF F by a block containing q cells of each of two colours. Clauses require exactly one cell to be selected, and its colour represents the value of the original variable. There is a family of projections which leave one cell of each colour active in every block and recover the base formula F. The lifted formula is still a CNF, and its size is polynomial when F and q have polynomial size. This style of lifting has been used in many other results in the literature, see e.g. [8, 51, 9, 33].

Consider a small depth-four refutation of this lifted formula. A bottom monomial which uses many variable blocks has only a small probability of surviving a random projection. Monomials which survive and involve few blocks become polynomials of small degree in the original variables. A union bound over all distinct bottom monomials therefore gives a projection under which every bottom polynomial has bounded degree, provided the original proof was sufficiently small. Since the projection recovers F and preserves semantic inferences, it produces a refutation to which the bounded-bottom-degree lower bound applies. Instantiating this construction with the random logarithmic-width CNFs gives the $n^{\Omega(\log n)}$ bound for unrestricted semantic $\Sigma\Pi\Sigma\Pi$-PC.

For the unary pigeonhole principle, we can simplify the proof by sparsifying its underlying graph directly. Independently retain each pigeon–hole edge with probability Δ/n, where $\Delta = \Theta(\ln n)$, and set the variables of the other edges to zero. With high probability, the remaining graph has the expansion properties needed for the rank argument. After multilinearising and working modulo the hole axioms, a non-zero bottom monomial of degree greater than w survives with probability at most $(\Delta/n)^{w+1}$. Taking a union bound over the distinct bottom monomials in a small proof gives a refutation of a sparse graph pigeonhole principle with bounded bottom degree. Balancing this survival estimate against the graph lower bound yields the superpolynomial bound for PHP_n^{2n}.

1.2.4 Depth-d $\mathsf{AC}^0[2]$-Frege Lower Bounds

The depth of the input formulas is an essential qualification for a semantic hierarchy result. A semantic system which can efficiently represent conjunctions of the input axioms can combine them in a balanced

tree and obtain a short refutation — this is already possible if we allow one more layer beyond $\Sigma\Pi\Sigma\Pi$. Thus, to obtain lower bounds at larger semantic proof depths, the depth of the input formulas must also be taken into account.

The proofs for depths three and four share a common strategy. At depth three, multiplication by a suitable polynomial annihilates the high-codimension product terms. For $\Sigma\Pi\Sigma\Pi^{[w]}$-PC, it annihilates the errors introduced by replacing these terms with low-degree approximations. It turns out that we can extend this latter argument to depth-d $\mathsf{AC}^0[2]$-Frege by approximating gates from the bottom up throughout the circuits. The same rank bounds will then allow us to find a common annihilator for the errors throughout a small proof, and multiplication by this annihilator makes the approximations exact, so semantic inferences remain sound.

The hardest part in the new argument is ensuring that the annihilator is hard to derive in depth-d $\mathsf{AC}^0[2]$-Frege. To put it in simple terms: when approximating a depth-d formula by the Razborov-Smolensky-style argument, we will necessarily introduce a more complicated set of errors that need to be eliminated by a common annihilating polynomial f when compared to a depth-4 formula. Naively, in order to cover this more complicated set of errors, we would be forced to use an annihilating polynomial f with larger degree. But, since we are ultimately reducing to a PC degree lower bound, it will be harder to guarantee that the existing lower bounds suffice.

To get around this issue, as in the unbounded bottom fan-in depth-four lifting argument, we start with a CNF requiring large PC degree and replace its variables by small gadgets. However, here the gadgets must produce instances of depth d after substitution, due to the above discussion. We will choose the gadgets so that a large space of candidate annihilators inherits hardness from the original CNF, in the same sense as in the depth-3 and depth-4 bounds: a derivation of a candidate annihilator f using lines fQ with Q of low degree can be transferred to a low-degree refutation of the original CNF. Crucially, our choice of gadget g will allow us to perform this transfer in a way that depends *only* on the degree of Q, *independently* of the degree of f. This lets us use a sufficiently large space of multipliers to annihilate all approximation errors while retaining the contradiction with the original degree lower bound. The gadget we choose will therefore have two roles: it raises the depth of the instance, and preserves enough of the original degree hardness for the annihilation argument.

Let F be the original CNF in variables $z_1, \ldots, z_n$. For each z_i, introduce a block

$$X_i := (x_{i,1}, \ldots, x_{i,m})$$

of fresh Boolean variables, with distinct blocks using distinct variables. Choose a small Boolean formula "gadget" g on m inputs, built from AND and OR gates with negations at the inputs. We use two gadgets on each block: $g_0(X_i) := g(X_i)$ and $g_1(X_i) := g(\overline{X_i})$, where $\overline{X_i} := (\neg x_{i,1}, \ldots, \neg x_{i,m})$. We replace the negative literal $\neg z_i$ by $g_0(X_i)$ and the positive literal z_i by $g_1(X_i)$. Thus $g_b(X_i) = 1$ asserts that the block encodes $z_i = b$.

We require the two circuits to have disjoint sets of satisfying assignments:

$$g_0(X) \wedge g_1(X) = 0 \quad \text{for every } X \in \{0,1\}^m.$$

In other words, no assignment to a block can make both opposite literals true. Both circuits are allowed to be false. Any assignment satisfying all substituted clauses therefore gives a satisfying assignment to F: set $z_i = b$ whenever $g_b(X_i) = 1$, and choose z_i arbitrarily when both circuits are false. Every substituted clause has a true literal, and its corresponding base literal remains true under this decoding. Hence the substitution preserves unsatisfiability.

To obtain axioms of the desired depth, we choose g with an AND root and depth at most $d + 1$. Distributing each substituted clause over these AND roots expresses it as a conjunction of depth-d circuits, which we take as the lifted axioms. Since the original clauses have constant width, this produces only polynomially many axioms when g has polynomial size.

For the degree transfer, we also require g to be *antitone*: changing an input from 0 to 1 cannot change its output from 0 to 1. Its family of accepted supports,

$$\mathcal{L} := \{S \subseteq [m] : g(1_S) = 1\},$$

is therefore downward closed, where 1_S denotes the indicator vector of S. For a degree cutoff k, we define the space of candidate annihilators by

$$U := \operatorname{span}\left\{ \prod_{i=1}^{n} \prod_{j \in S_i} x_{i,j} : S_i \in \mathcal{L} \text{ for every } i \in [n],\ \sum_{i=1}^{n} |S_i| \le k \right\}.$$

Thus the monomials spanning U have total degree at most k, and their support in each block is accepted by g.

Antitonicity gives us a key algebraic property: the monomials $\prod_{j \in S} x_j$ with $S \in \mathcal{L}$ form a basis for functions on either of the two sets $\{X : g_0(X) = 1\}$ and $\{X : g_1(X) = 1\}$. Taking products of these monomials across blocks gives a common basis on the encodings of every base assignment z. In this basis, each $f \in U$ has the same constant coefficients for every z. Multiplying f by a degree-t monomial changes at most t blocks, so the resulting coefficients can depend on at most t base variables. By linearity, the coefficients of fQ are therefore polynomials in the base (unlifted) variables $z_1, \dots, z_n$ of degree at most $\deg Q$, independently of $\deg f$.

As we prove in Section 8.2, extracting these coefficients transfers a PC derivation of f, with every line of the form fQ, from the lifted axioms multiplied by f to a PC derivation from F. At the final line, a non-zero $f \in U$ has a coefficient equal to 1, so the transferred derivation is a refutation. Its degree is controlled by the degrees of the factors Q, even when f itself has large degree.

As in the preceding arguments, we need $\dim(U)$ to exceed the total rank of the approximation errors. A natural choice is the complement of the majority function, $\overline{\mathrm{Maj}}$, which has the required properties and accepts enough monomial supports to make U large. However, a direct substitution of $\overline{\mathrm{Maj}}$ would not work for our construction, since $\overline{\mathrm{Maj}}$ does not have polynomial-size depth-d $\mathsf{AC}^0[2]$ circuits. So, instead, we will use the coin distinguishers of Limaye et al. [52] to construct a small approximate majority gadget, which will retain enough of the properties of majority needed while keeping $\dim(U)$ large.

Combining this lifting with the approximation argument gives the superpolynomial lower bound in Theorem 8.1. One additional layer permits conjunction of the lifted axioms, yielding a polynomial-size semantic depth-$(d+1)$ AC^0-Frege refutation. Thus the same instances separate consecutive semantic proof depths, even though the lower-depth system also permits parity gates.

1.3 Comparison with concurrent results

Pang's results. In independent recent work, Pang [55] proves a degree–size relation for resolution over bounded-degree polynomial equations over prime fields, denoted $\mathsf{Res}(\mathsf{PC})$, with applications to random CNFs and several pigeonhole principles. Building on Braun's work, his argument also shows how ordinary polynomial calculus degree lower bounds can be used to obtain size lower bounds for syntactic $\mathsf{Res}(\mathsf{PC}_w)$. In particular, he obtains a $2^{\Omega(n)}$ lower bound for $\mathsf{Res}(\oplus)$ refutations of the bit pigeonhole principle with $n+1$ pigeons and n holes.

Exponential lower bounds for $\mathsf{Res}(\mathsf{PC}_w)$ over $\mathbb{F}_2$ in the range $w \le \varepsilon \log n$ were also obtained independently in the present work. Our bounds for random CNFs of logarithmic width give $\exp\left(n^{\Omega(1)}\right)$ line lower bounds even for semantic $\mathsf{Res}(\mathsf{PC}_w)$, by Theorems 2.4 and 7.2. This overlaps with Pang's degree range, but does not give his general degree–size relation from a PC degree lower bound alone.

In general, some of our results are stronger than Pang's, while others are not directly comparable to his. The main difference is that we work with semantic proof systems, whereas Pang works with syntactic ones. This means in particular that our lower bounds apply to stronger proof systems, up to full semantic depth-d $\mathsf{AC}^0[2]$-Frege, for every fixed d, on depth-d instances. Pang, in turn, exploits the syntactic derivation rules to obtain a general degree-size tradeoff that does not extend to the corresponding semantic systems.

Our results strengthen his $\mathsf{Res}(\oplus)$ consequences. For the same bit pigeonhole principle, we prove a $2^{\Omega(n)}$ lower bound on the number of distinct non-zero top product terms in semantic $\Sigma\Pi\Sigma$-PC (Corollary 6.14). Thus the exponential bound persists when proof lines may be arbitrary sums of products of affine forms and every sound binary inference is allowed. The same extension holds for the functional pigeonhole principle and random 3-CNFs (Corollary 6.13 and Theorem 6.5).

Over $\mathbb{F}_2$, semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC simulates syntactic $\mathsf{Res}(\mathsf{PC}_w)$ line by line, with linear overhead in total algebraic size (Theorem 2.4). The simulation replaces each disjunction of polynomial equations by the product of its polynomials. Our system additionally allows sums of these products and arbitrary sound binary inferences. Thus our lower bounds also apply to $\mathsf{Res}(\mathsf{PC}_w)$ on the same instances.

At depth four, we prove superpolynomial total circuit-size lower bounds with no restriction on bottom degree. These hold for the weak unary pigeonhole principle and for lifted random CNFs (Theorems 7.12 and 7.14). By the simulation above, they also give lower bounds on total circuit size for unrestricted-degree $\mathsf{Res}(\mathsf{PC})$ on the same instances. These results go beyond the bounded-degree setting of Pang's lower bounds.

For every fixed $d \geq 2$, we also construct depth-d instances requiring superpolynomial-size semantic depth-d $\mathsf{AC}^0[2]$-Frege refutations, while admitting polynomial-size refutations one depth higher (Theorem 8.1). This is an unconditional lower bound at every fixed proof depth.

There is also a difference in the hypotheses needed for the lower bounds. Pang derives his degree–size tradeoff from a PC degree lower bound alone. Our algebraic applications also require quantitative lower bounds on the *ranks of pseudo-reduction operators*. Both proofs use rank estimates to choose a common multiplier; the distinction is in the lower-bound information required from the input formula. Pang's results hold over arbitrary prime fields, whereas our results are over $\mathbb{F}_2$.

The methods are closely related. Both use polynomial approximation and a common multiplier f. Pang requires the approximation errors multiplied by f to have bounded-degree PC derivations from the input axioms. He then uses the specified syntactic rules to derive the multiplied conclusion of each inference [55, Proposition 5.1]. We instead make the approximation errors vanish after multiplication by f. The resulting polynomials agree with the original proof lines wherever f is non-zero. This way, arbitrary sound binary inferences remain valid. Our parallel syntactic paper [4] uses majority substitution to extend the degree range beyond these independently obtained logarithmic-degree lower bounds.

The input families also matter. Pang's (syntactic) degree–size tradeoff applies to constant-width CNFs whose semantic depth-four refutations are short when the bottom degree is at least the clause width. Consequently, our stronger systems require different hard instances in that degree range.

Conneryd's results. A second closely related work which appeared during the preparation of our paper is that of Conneryd [24]. For $k = O(\log n)$, he proves that a random k-CNF on n variables with $n2^k$ clauses requires PC degree at least $n/2 - O(n/\sqrt{k})$ over every field, with high probability. Applying Pang's $\mathsf{AC}^0[p]$-Frege tradeoff, he obtains a depth-two $\mathsf{AC}^0[p]$-Frege lower bound

$$\exp\Big(\Omega_p(\sqrt{k}\log n)\Big)$$

for every fixed prime p. In particular, $k = \Theta(\log n)$ gives a lower bound of $\exp\Big(\Omega_p\Big((\log n)^{3/2}\Big)\Big)$.

Here depth two is the depth of each entire Boolean proof line over unbounded-fan-in AND, OR and modular p gates; negations do not contribute to depth. Thus both CNF and DNF lines are allowed. This is the usual syntactic Frege system, with any fixed finite set of sound, implicationally complete rule schemes and the standard modular axioms [24, Section 2.3]. This uses the Frege tradeoff in Section 8 of [55]. Conneryd's Appendix A extends that tradeoff to these rule schemes. The syntactic rules are used to derive the inference errors from local approximation errors in bounded PC degree [24, Lemma A.8]. In the semantic version, every unsatisfiable CNF has a short depth-two refutation by successively conjoining its clauses. Our semantic $\mathsf{AC}^0[2]$-Frege lower bounds therefore concern deep instances instead.

Conneryd also proves a nearly matching PC degree upper bound of

$$\frac{n}{2} - \frac{n\ln k}{2k} + O\Big(\frac{n}{k}\Big).$$

This is an upper bound on degree, rather than on Frege proof size. It shows that these random formulas do not satisfy the degree hypothesis of the stated tradeoff for Frege depth at least three [24, Theorem 1.3].

Comparison with our parallel syntactic paper. In a parallel syntactic paper (with Shekhovtsov) [4] we extend these independently obtained lower bounds for $\mathsf{Res}(\mathsf{PC}_w)$ to larger degrees by majority substitution. A line in $\mathsf{Res}(\mathsf{PC}_w)$ has the form $L = \bigvee_{i=1}^{m}(u_i = 1)$, where each u_i is a polynomial over $\mathbb{F}_2$ of degree at most w. Its Boolean form is

$$\mathsf{OR} \circ \oplus \circ \mathsf{AND}_{\leq w}.$$

The paper also proves superpolynomial lower bounds for its syntactic depth-3.5 $\mathsf{AC}^0[2]$-Frege system. In its refutation presentation, a line has the form $L = \bigvee_{i=1}^{m} \bigwedge_{j=1}^{t_i}(u_{i,j} = 1)$, where each $u_{i,j}$ is a polynomial over $\mathbb{F}_2$ of degree at most b. Each such polynomial is a parity of monomials of degree at most b. Thus the Boolean form of a line is

$$\mathsf{OR} \circ \mathsf{AND} \circ \oplus \circ \mathsf{AND}_{\leq b}.$$

The top OR, AND and parity gates have unbounded fan-in. The bottom AND gates have fan-in at most b. The dual presentation in the introduction of [4] has the form $\mathsf{AND} \circ \mathsf{OR} \circ \oplus \circ \mathsf{AND}_{\leq b}$.

The construction gives CNF families of size N whose refutations require superpolynomial size when $b = o(\sqrt{\log N})$. The syntactic restriction is essential for this class of proof lines. Already for $b = 1$, the permitted proof lines include arbitrary DNFs, so the semantic version has short refutations of every unsatisfiable CNF.

Both papers use polynomial approximation, a common multiplier and lifting. In the syntactic paper, the multiplied error witnesses need only be consequences of the input axioms in bounded PC degree. The resulting derivation follows the specified syntactic rules.

In the current paper we choose a non-zero common multiplier f so that, for every original proof-line polynomial P and its low-degree approximation $\widetilde{P}$,

$$f\widetilde{P} = fP.$$

Equivalently, multiplication by f annihilates the approximation error: $f(\widetilde{P} - P) = 0$. Consider a Boolean assignment a satisfying the "background axioms" (e.g., the hole axioms in the PHP), if any. If $f(a) = 1$, the identity gives $\widetilde{P}(a) = P(a)$, so the approximation equals the original line on this assignment. If $f(a) = 0$, both multiplied polynomials are zero. Thus, on each such assignment, the multiplied inference either agrees with the original inference or has all its premise and conclusion polynomials equal to zero. This preserves every sound (binary) inference.

Both papers also obtain total-size lower bounds for unrestricted-degree $\mathsf{Res}(\mathsf{PC})$. On the instances constructed here, these bounds hold for semantic $\Sigma\Pi\Sigma\Pi$-PC, which additionally permits sums of the products encoding $\mathsf{Res}(\mathsf{PC})$ lines and arbitrary sound binary inferences (Theorem 2.4). For $\mathsf{AC}^0[2]$-Frege, the syntactic paper proves CNF lower bounds with the specified line structure and restricted bottom degree. Here we prove semantic lower bounds at every fixed depth for deep instances.

Semantic versus syntactic lower bounds and limitations Both the semantic and syntactic approaches to $\mathsf{AC}^0[2]$-Frege lower bounds for CNFs have limitations. Pang's syntactic approach gives a sufficient condition for CNF $\mathsf{AC}^0[2]$-Frege lower bounds [55, Proposition 8.2]. Fix $d \geq 2$ and a family of polynomial-size unsatisfiable CNFs F_n on n variables. Write $w(F_n)$ for the maximum clause width. Suppose that the polynomial encoding of F_n has no PC refutation over $\mathbb{F}_2$ of degree at most D, where $D - w(F_n) \geq \frac{n}{2} - o\left(\frac{n}{(\log n)^{d-2}}\right)$. Then every syntactic depth-d $\mathsf{AC}^0[2]$-Frege refutation of F_n has size $n^{\omega(1)}$. Establishing this condition for every fixed d remains open, and is a form of SETH (Strong Exponential Time Hypothesis) for PC.

On the other hand, our semantic approach applies to stronger proof systems, but has a different limitation. Already at depth five, semantic $\Sigma\Pi\Sigma\Pi\Sigma$-PC has polynomial-size refutations of every unsatisfiable CNF. Thus extending our argument to CNF $\mathsf{AC}^0[2]$-Frege lower bounds requires using restrictions on the syntactic inference rules. Our semantic $\mathsf{AC}^0[2]$-Frege lower bounds instead use instances of the maximum allowed depth.

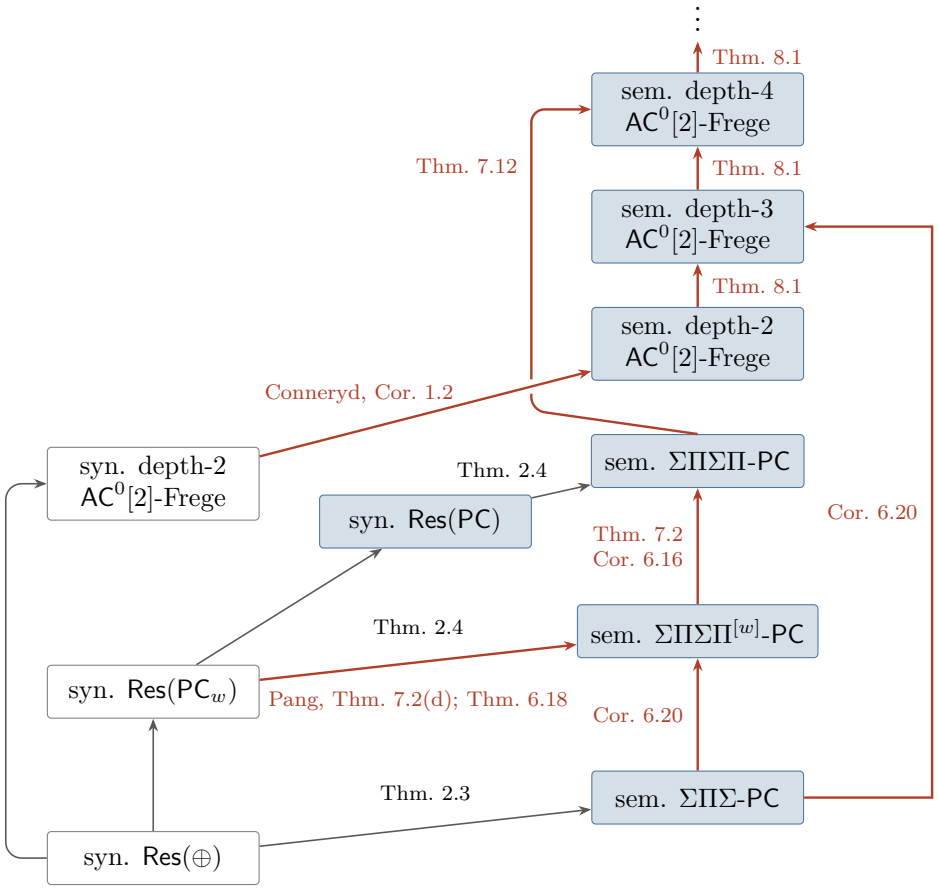


Figure 1: Simulations and separations over $\mathbb{F}_2$, for fixed $w \geq 2$. Blue boxes mark systems for which we prove lower bounds in this paper; white boxes are comparison systems. An arrow $A \to B$ means that B polynomially simulates A; red arrows indicate strict simulations. The labels "syn." and "sem." denote syntactic and binary semantic inference, respectively.

1.4 Paper Outline

The rest of the paper is outlined as follows. Section 2 introduces the proof systems and establishes the simulations of Res(⊕) and Res(PC). Section 3 gives a re-proof of Braun's result [16], using the annihilating polynomial method, for a lower bound for Res(⊕) proofs of the bit pigeonhole principle. Section 4 generalises this, and provides an abstract lower bound framework using annihilating polynomials, low-degree approximations, and pseudo-reduction operators. Section 5 proves the rank estimates for expansion pseudo-reduction operators needed in our applications. Section 6 proves exponential semantic ΣΠΣ-PC lower bounds for random 3-CNFs and the functional pigeonhole principle, yielding separations from semantic ΣΠΣΠ-PC. Section 7 establishes the semantic ΣΠΣΠ-PC lower bounds, removing bottom-degree restrictions by selector lifting and sparsification. Section 8 proves the semantic AC^0[2]-Frege depth hierarchy theorem by lifting CNFs with large PC refutation degree.

2 Preliminaries

Linear-algebra notation. Unless stated otherwise, all vector spaces and linear maps are over the field currently under consideration, and all vector spaces are finite-dimensional.

For a subset S of a vector space, $\operatorname{span}(S)$ denotes its linear span, and $\dim(V)$ denotes the dimension of a vector space V. If U is a linear subspace of V, we write

$$\operatorname{codim}_V(U) = \dim(V) - \dim(U),$$

and omit the subscript when the ambient space is clear. For a linear map $R\colon V \to W$, we use $\ker(R)$ and $\operatorname{im}(R)$ for its kernel and image, respectively, and set $\operatorname{rank}(R) = \dim(\operatorname{im}(R))$. We write $V = U \oplus W$

if every vector in V has a unique representation as a sum of a vector in U and a vector in W; more generally, $\bigoplus_{i\in I} V_i$ denotes the direct sum of the spaces V_i (equivalently, their direct product when I is finite).

Algebraic circuits. We consider algebraic circuits computing polynomials over a field $\mathbb{F}$. An algebraic circuit is a labelled directed acyclic graph. The leaves of the graph (nodes of in-degree zero) are labelled with variables or field elements. They *compute* the polynomials on their label (either a variable or a field constant). There are two types of internal nodes: plus and times gates. A plus gate (denoted by Σ) computes an arbitrary linear combination of the polynomials computed at its input nodes, with the coefficients specified by the weights of the corresponding incoming edges. A product-gate (denoted Π) computes the product of the polynomials computed in its incoming nodes. Accordingly, the fan-in of internal gates is unbounded. The polynomial computed by the algebraic circuit is the polynomial computed in the (single) output gate (node of zero outgoing degree).

We consider layered circuits with alternating Σ- and Π-layers, numbered from the output towards the inputs. The *depth* of a (layered) circuit is the maximal number of layers from the output gate toward the input layer. The *size* $|C|$ of an algebraic circuit C is the total number of vertices in its directed acyclic graph. Field elements and edge weights have unit cost.

For bounded-depth circuits, we use the following notation:

- A $\Sigma\Pi$ circuit is a depth-two circuit. Equivalently, it represents a polynomial as a linear combination of monomials.
- A $\Sigma\Pi\Sigma$ circuit is a depth-three circuit. Because a bottom Σ-gate may receive a field constant, its output is an affine form $a_0 + \sum_i a_i x_i$. We call the polynomials computed by its second-layer Π-gates *product terms*. Thus, every product term is a product of affine forms.
- A $\Sigma\Pi\Sigma\Pi$ circuit is a depth-four circuit. We write $\Sigma\Pi\Sigma\Pi^{[w]}$ for the subclass in which every bottom-layer Π-gate has fan-in at most w. We call the polynomials computed by the Π-gates in the second (from top) layer *top product terms*.

2.1 Propositional proof systems

Propositional proofs certify the unsatisfiability of Boolean formulas, normally in CNF. Typically, a proof system operates with proof lines that are more expressive than the clauses of the original CNF. Its inference rules allow one to derive new proof lines from previously derived ones, and a refutation consists of deriving an identically false proof line from the clauses of the formula, such as $1 = 0$.

For example, in the Polynomial Calculus proof system, proof lines are polynomials, or more precisely, equations asserting that a polynomial is equal to zero.

We identify each literal with its truth-value polynomial: $1 - x$ for a positive literal x and x for a negative literal $\neg x$. The polynomial representation of a clause

$$C = \bigvee_{i\in P} x_i \vee \bigvee_{i\in N} \neg x_i$$

is

$$[\![C]\!] = \prod_{i\in P} (1 - x_i) \cdot \prod_{i\in N} x_i$$

Thus, $[\![C]\!]$ is the indicator function of Boolean assignments that falsify C, and the clause is encoded by the polynomial $[\![C]\!]$: $C(x) = 1$ holds exactly when $[\![C]\!](x) = 0$.

Polynomial Calculus. Here we define PC [23], which we sometimes call "syntactic PC" to distinguish it from semantic PC. Let $\mathbb{F}$ be a field and let

$$\mathcal{F} = \{f_1, \ldots, f_m\} \subseteq \mathbb{F}[x_1, \ldots, x_n]$$

represent the system of equations

$$f_1 = \cdots = f_m = 0.$$

Our goal is to show that this system has no solution on the Boolean cube $\{0,1\}^n$.

For Boolean variables, the identities

$$x_i^2 - x_i = 0$$

hold for all $i \in [n]$. Therefore, we will consider polynomials modulo the ideal generated by all Boolean identities, that is, we will work in the quotient ring

$$\mathcal{B}_n(\mathbb{F}) = \mathbb{F}[x_1, \ldots, x_n]/\langle x_i^2 - x_i : i \in [n]\rangle.$$

Thus, a polynomial or circuit in a proof line represents its residue class in $\mathcal{B}_n(\mathbb{F})$ (in particular, a circuit need not be expanded or multilinearised when written as a proof line). All equalities in the inference rules are interpreted in this quotient. Circuit size refers to the displayed representation, whereas degree refers to its unique multilinear representative.

Every element of $\mathcal{B}_n(\mathbb{F})$ has a unique multilinear representative. Moreover, evaluation on the Boolean cube identifies $\mathcal{B}_n(\mathbb{F})$ with the space of all functions $\{0,1\}^n \to \mathbb{F}$. In particular, two elements of $\mathcal{B}_n(\mathbb{F})$ are equal if and only if they agree on $\{0,1\}^n$.

This equivalence is important for us, and for this reason, if $p, q \in \mathbb{F}[x_1, \ldots, x_n]$ then we write

$$p \equiv q \iff p = q \bmod \langle x_i^2 - x_i\rangle_{i=1}^n \iff \forall x \in \{0,1\}^n : p(x) = q(x).$$

A Polynomial Calculus (PC) [23] derivation from $\mathcal{F}$ is a sequence of polynomials in which every line is an axiom or follows from earlier lines by one of the rules

$$\frac{p \qquad q}{\alpha p + \beta q} \quad (\alpha, \beta \in \mathbb{F}), \qquad \frac{p}{x_i p} \quad (i \in [n]).$$

A PC refutation is a derivation of 1 (standing for $1 = 0$), which certifies that the equations have no common Boolean solution. The *degree* of a PC derivation $\pi = (p_1, \ldots, p_L)$ is $\max_{j\in[L]} \deg(p_j)$.

The *size* of a refutation or a derivation depends on how the polynomials occurring in them are represented. We are concerned with the following measures:

1. *Sum of monomials representation*: each polynomial is written explicitly as a sum of monomials with coefficients from the field. The size is defined as the total number of monomials in all the proof lines. Equivalently, this representation can be viewed as a $\Sigma\Pi$ algebraic circuit: a depth-2 arithmetic circuit whose top gate computes an affine linear combination, while the gates at the next level compute products of variables, which form the inputs to the circuit.

2. *Monomial size*: a lower bound on the above Item 1 is given by the *monomial size* of the proof, defined as the total number of *distinct* monomials that occur in its proof lines.

3. *Algebraic circuit size*: If $\mathcal{C}$ is a class of algebraic circuits, then $\mathcal{C}$-PC denotes Polynomial Calculus in which every line is represented by a circuit from $\mathcal{C}$. The $\mathcal{C}$-size of a proof π, denoted size(π), is the number of distinct rooted subcircuits appearing anywhere in π, allowing sharing across proof lines.

4. *Top product term*: Semantic systems use the same representation-dependent size measures as their syntactic counterparts. Our lower bounds concern $S(\pi)$ (see below), the number of distinct non-zero top product terms across the derivation, identified as elements of the Boolean quotient. This quantity is a lower bound on total circuit size.

Polynomial Calculus over bounded-depth arithmetic circuits. We consider $\mathcal{C}$-PC for

$$\mathcal{C} \in \left\{\Sigma\Pi\Sigma, \Sigma\Pi\Sigma\Pi, \Sigma\Pi\Sigma\Pi^{[w]}\right\}.$$

For a derivation π, let $\mathcal{T}(\pi)$ denote the set of all distinct non-zero top product terms appearing in π, viewed as elements of $\mathcal{B}_n(\mathbb{F})$. We use

$$S(\pi) = |\mathcal{T}(\pi)|$$

as a complexity measure. Clearly, $S(\pi)$ is a lower bound on the $\mathcal{C}$-size of π.

Raz and Tzameret [57, Section 9] (cf. [58]) used the clause-to-product translation to simulate (a fragment of) resolution over linear equations by PC over depth-three (multilinear) circuits with a sum gate at the root. Over $\mathbb{F}_2$, the same translation gives our line-wise simulation of semantic Res($\oplus$) by semantic $\Sigma\Pi\Sigma$-PC, with one product term per clause (Theorem 2.3). Moreover, Impagliazzo, Mouli, and Pitassi [39] use affine extension variables in $\Sigma\Pi\Sigma$-PC; substituting their definitions yields sums of products of affine forms, as in our proof lines. Similarly, Grigoriev and Hirsch [34] considered algebraic proof systems over algebraic formulas as well as constant-depth variants.

Resolution over parities. Resolution over parities (Res($\oplus$)) was studied by Itsykson and Sokolov [45] as the mod-2 variant of resolution over linear equations by Raz and Tzameret [57]. A proof line in this system is a *linear clause*, defined as a disjunction

$$C = (\ell_1 = b_1) \vee \cdots \vee (\ell_w = b_w),$$

where each $\ell_i = \sum_{j=1}^{n} a_{i,j} x_j$ is a linear form over $\mathbb{F}_2$ and $b_i \in \mathbb{F}_2$. Ordinary clauses are included by writing the literals x_i and $\neg x_i$ as the equations $x_i = 1$ and $x_i = 0$, respectively.

A Res($\oplus$) derivation from a CNF formula consists of linear clauses. Each line is an axiom clause or follows from earlier lines by resolution or weakening:

$$\frac{A \vee (\ell = 0) \qquad B \vee (\ell = 1)}{A \vee B}, \qquad \frac{C}{D} \quad \text{(whenever } C \models D\text{)},$$

where ℓ is a linear form and $C \models D$ means that every Boolean assignment satisfying C also satisfies D. A refutation derives the empty clause. We allow lines to be reused, so proofs are DAG-like, and define the size $|\pi|$ to be the number of lines, including axiom lines.

2.2 Semantic proof systems

Semantic proof systems retain the permitted form of proof lines, but dispense with particular inference rules. Identify a line A with the set $[A]$ of Boolean assignments that satisfy it. A sound binary inference from A and B to C is then exactly the containment

$$[A] \cap [B] \subseteq [C].$$

The semantic system permits every such inference for which C is again an allowed proof line. It is therefore the strongest sound binary system with the given class of lines: its complexity reflects what one line can express, rather than the limitations of a particular set of rules [49, Section 17.5].

Semantic systems are not assumed to admit polynomial-time proof verification, hence they are not necessarily Cook–Reckhow proof systems [25]. This is not a restriction, since we are concerned in this work with *lower bounds*. Note that semantic resolution over parities does admit polynomial-time verification: a binary inference between linear clauses can be checked by Gaussian elimination.

This relaxation is useful for lower bounds. Every sound syntactic inference of arity at most 2 is also a semantic inference, so a lower bound for the semantic system applies immediately to its syntactic counterpart. Bounding the arity is essential: with unbounded arity, all axioms of an unsatisfiable instance would imply a contradiction in a single inference. We use arity 2 throughout.

Semantic Polynomial Calculus. Let $\mathbb{F}$ be any field and let $\mathcal{C}$ be a class of algebraic circuits over $\mathbb{F}$. Arity-2 semantic $\mathcal{C}$-PC has the same axioms as $\mathcal{C}$-PC, but replaces the syntactic inference rules by

$$\frac{P \qquad Q}{R} \qquad \text{whenever} \qquad P(x) = Q(x) = 0 \implies R(x) = 0 \quad \text{for every } x \in \{0,1\}^n,$$

where $P, Q, R \in \mathcal{C}$. Thus, any polynomial that is a semantic consequence of two previously derived polynomials may be derived. As usual, a refutation ends with the constant polynomial 1, and the degree is the maximum degree of the multilinear polynomials represented by its lines.

Over $\mathbb{F}_2$, every arity-2 semantic inference can be simulated in ordinary PC with only a constant-factor increase in degree. The following is the case $p = 2$ and $k = 2$ of [19, Lemma 15].

Lemma 2.1 (PC simulates semantic PC with only a factor of three increase in degree)**.** *Let $f_1, f_2, g \in \mathcal{B}_n(\mathbb{F}_2)$ have degree at most d, and suppose that $f_1(x) = f_2(x) = 0$ implies $g(x) = 0$ for every $x \in \mathbb{F}_2^n$. Then there are multilinear polynomials $h_1, h_2 \in \mathcal{B}_n(\mathbb{F}_2)$ of degree at most $2d$ such that*

$$g \equiv h_1 f_1 + h_2 f_2.$$

Consequently, g has a syntactic PC *derivation from f_1 and f_2 of degree at most $3d$.*

Proof. By soundness, g vanishes whenever $f_1 = f_2 = 0$. Since we work over $\mathbb{F}_2$, this implies

$$g(1 + f_1)(1 + f_2) \equiv 0,$$

and therefore

$$g \equiv gf_1 + gf_2 + gf_1f_2 = g(1 + f_2)f_1 + gf_2.$$

Thus we may take $h_1 = g(1 + f_2)$ and $h_2 = g$, of degree at most $2d$ and d, respectively. Expand g into monomials. Multiplying f_1 and f_2 by these monomials, one variable at a time, and adding the results derives gf_1 and gf_2 in degree at most $2d$. Likewise, multiplying gf_1 by the monomials of f_2 and adding derives gf_1f_2 in degree at most $3d$. Adding these three polynomials gives g. All equalities are in $\mathcal{B}_n(\mathbb{F}_2)$, so multilinearisation does not increase degree. □

Applying the lemma to each inference shows that every degree-d arity-2 semantic $\mathcal{C}$-PC derivation over $\mathbb{F}_2$ can be transformed into a syntactic PC derivation of degree at most $3d$, from the same axioms and with the same conclusion. An inference with one premise is covered by taking $f_1 = f_2$.

Semantic resolution over parities. Arity-2 semantic Res($\oplus$) has the same axioms as Res($\oplus$) and replaces its inference rules by

$$\frac{C \qquad D}{E} \qquad \text{whenever} \qquad C \wedge D \models E,$$

for arbitrary linear clauses C, D, and E. Thus every Boolean assignment that satisfies both premises must satisfy the conclusion. This rule subsumes weakening by taking $C = D$.

Itsykson and Sokolov [45] showed that every semantic Res($\oplus$) derivation of size S can be simulated by a syntactic Res($\oplus$) derivation of size at most $3S$.

Semantic Res(PC) **over a field $\mathbb{F}$.** In semantic Res(PC), proof lines are finite disjunctions of polynomial equations

$$(f_1 = 0) \vee (f_2 = 0) \vee \cdots \vee (f_k = 0),$$

where $f_1, \ldots, f_k \in \mathbb{F}[x_1, \ldots, x_n]$. The system permits any semantically sound inference from at most two premises: a conclusion may be derived whenever every Boolean assignment satisfying the premises also satisfies the conclusion. The corresponding syntactic proof system was introduced over $\mathbb{F}_2$ by Krajíček [50] and subsequently studied over finite rings by Khaniki [47]. We write Res(PC_d) for the subsystem of Res(PC) in which every polynomial appearing in a proof line has degree at most d.

Each polynomial is represented explicitly as a sum of monomials, viewed as a $\Sigma\Pi$ circuit. The size size(π) counts distinct rooted subcircuits across the entire derivation, allowing sharing across proof lines. This extends the monomial size measure for PC, which counts distinct monomials across all proof lines, by also counting the gates combining them.

Semantic $\mathsf{AC}^0[2]$-Frege Our $\mathsf{AC}^0[2]$-circuits have Boolean variables, constants, and unbounded-fan-in AND, OR, and XOR gates, with negations only at inputs. Consecutive equal gate types are flattened; repeated children are removed for AND/OR and cancelled in pairs for XOR. Empty AND gates equal 1, and empty OR/XOR gates equal 0. Depth counts gates; literals and constants have depth zero. Complementation preserves depth by De Morgan's laws and by adding a constant 1 input to an XOR gate.

The arithmetisation $[\![A]\!]$ of an $\mathsf{AC}^0[2]$ circuit A extends the polynomial translation $[\![C]\!]$ for clauses. We encode the *negation* of A as a polynomial $[\![A]\!]$ such that $A(x) = 1$ exactly when $[\![A]\!](x) = 0$. For constants and literals, $[\![0]\!] = 1$, $[\![1]\!] = 0$, $[\![x_i]\!] = 1 + x_i$, and $[\![\neg x_i]\!] = x_i$. For gates,

$$[\![\bigvee_i A_i]\!] = \prod_i [\![A_i]\!],$$
$$[\![\bigwedge_i A_i]\!] = 1 + \prod_i (1 + [\![A_i]\!]),$$
$$[\![\bigoplus_i A_i]\!] = 1 + \sum_i (1 + [\![A_i]\!]).$$

All operations are in $\mathcal{B}_n$, with empty products equal to 1 and empty sums to 0.

A depth-d semantic $\mathsf{AC}^0[2]$-Frege proof from a family of circuits $\mathcal{F}$ is a sequence of circuits of depth at most d. Each line is either an axiom from $\mathcal{F}$, a Boolean tautology, or a semantic consequence of at most two earlier lines. A refutation ends with the constant-zero circuit. Our primary complexity measure is *proof size*: given an $\mathsf{AC}^0[2]$-Frege proof π, its size $\mathrm{size}(\pi)$ is the number of distinct rooted subcircuits appearing anywhere in π. Observe that the measure $S(\pi)$ defined above is similar to this one: it counts the number of distinct top product gates in a proof. $\mathrm{size}(\pi)$ is more general and counts the size of π considered as a (multi-output) Boolean circuit.

A stronger complexity measure is *proof length*: given an $\mathsf{AC}^0[2]$-Frege proof π, its length $\mathrm{len}(\pi)$ is the number of proof lines in π. For *syntactic* Frege systems, these measures are polynomially equivalent when the size of the initial axioms is taken into account.

Proposition 2.2 (cf. Lemma 4.4.6 in [48])**.** *Fix a Frege system. There is a constant $C > 0$, depending only on the system, such that for every unsatisfiable family of circuits $\mathcal{F}$,*

$$\mathrm{len}(\mathcal{F} \vdash) \le \mathrm{size}(\mathcal{F} \vdash) \le C\,\mathrm{len}(\mathcal{F} \vdash) + \mathrm{size}(\mathcal{F}).$$

Here $\mathrm{len}(\mathcal{F} \vdash)$ *and* $\mathrm{size}(\mathcal{F} \vdash)$ *denote the minimum length and minimum size, respectively, of a refutation of $\mathcal{F}$ in the fixed system, and* $\mathrm{size}(\mathcal{F})$ *is the number of distinct rooted subcircuits appearing in $\mathcal{F}$. The same bounds hold for depth-d refutations.*

Although depth preservation is not stated in [48, Lemma 4.4.6], it follows from the proof.

For semantic proof systems, these measures need not be polynomially equivalent: a single inference may introduce a line containing arbitrarily many new subcircuits. Thus size is the more natural complexity measure in this setting.

2.3 Simulation of $\mathsf{Res}(\oplus)$

We show that semantic $\Sigma\Pi\Sigma$-PC over $\mathbb{F}_2$ polynomially simulates semantic $\mathsf{Res}(\oplus)$. Given a CNF formula ϕ, let

$$\mathcal{F}_\phi = \{[\![C]\!] : C \text{ is a clause of } \phi\}$$

be its polynomial encoding.

Theorem 2.3 (Simulation of $\mathsf{Res}(\oplus)$)**.** *Suppose that a CNF formula ϕ on $n \ge 1$ variables has a semantic $\mathsf{Res}(\oplus)$ refutation with L lines. Then $\mathcal{F}_\phi$ has a semantic $\Sigma\Pi\Sigma$-PC refutation π' with $S(\pi') \le L$, where $S(\pi')$ counts its distinct non-zero product terms. Moreover, π' has $\Sigma\Pi\Sigma$-size $O(nL)$.*

Proof. To every linear clause $C = (\ell_1 = b_1) \vee \cdots \vee (\ell_t = b_t)$ we associate the polynomial

$$[\![C]\!] = \prod_{j=1}^{t} (\ell_j + b_j)$$

(recall we interpret this product in $\mathcal{B}_n(\mathbb{F}_2)$, following our convention for circuit representations).

Let $\pi = (C_1, \ldots, C_L)$ be a semantic $\mathsf{Res}(\oplus)$ refutation of ϕ. For every linear clause C and every Boolean assignment a, the product defining $[\![C]\!](a)$ is zero exactly when at least one equation in C is satisfied by a. Thus, if $C \wedge D \models E$, then

$$[\![C]\!](a) = [\![D]\!](a) = 0 \quad \Longrightarrow \quad [\![E]\!](a) = 0$$

for every Boolean assignment a. This is precisely the condition for a valid arity-2 semantic $\Sigma\Pi\Sigma$-PC inference. Initial clauses translate into their polynomial encodings, so the line-wise translation

$$\pi' = (\llbracket C_1 \rrbracket, \ldots, \llbracket C_L \rrbracket)$$

is a semantic $\Sigma\Pi\Sigma$-PC refutation of $\mathcal{F}_\phi$. The final empty clause translates into the empty product, $\llbracket \varnothing \rrbracket = 1$, as required.

Since every linear clause $C = (\ell_1 = b_1) \vee (\ell_2 = b_2) \vee \cdots \vee (\ell_t = b_t)$ is the negation of the system of linear equations $\bigwedge_{i=1}^{t}(\ell_i = b_i + 1)$, every non-tautological linear clause can be rewritten in an equivalent form in which the equations of the corresponding linear system are linearly independent. In particular, we may assume that every linear clause contains at most n equations, where n is the number of variables.

This implies that every non-zero $\llbracket C_i \rrbracket$ has circuit size $O(n)$. Since each line contains at most one non-zero product term, $S(\pi') \leq L$; summing the circuit sizes gives $O(nL)$. □

2.4 Simulation of Res(PC)

In this subsection, we show that, over an arbitrary field $\mathbb{F}$, semantic Res(PC) is simulated line-by-line by semantic $\Sigma\Pi\Sigma\Pi$-PC. Moreover, the simulation maps semantic Res(PC_w) into semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC.

Theorem 2.4 (Simulation of Res(PC))**.** *Let $\mathbb{F}$ be any field, and let φ be a CNF formula. Suppose that φ has a semantic* Res(PC) *refutation π, in which every polynomial is represented explicitly as a sum of monomials. Then the polynomial encoding $\mathcal{F}_\varphi$ has a semantic $\Sigma\Pi\Sigma\Pi$-PC refutation π' whose total circuit size is*

$$O(\text{size}(\pi)).$$

The simulation preserves the inference arity. Moreover, if π has L lines, then

$$S(\pi') \leq L.$$

If π is a semantic Res(PC_w) *refutation, then π' is a semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutation.*

Proof. Consider a proof line

$$D = (f_1 = 0) \vee \cdots \vee (f_k = 0)$$

of π, and associate with it the polynomial

$$\llbracket D \rrbracket = \prod_{i=1}^{k} f_i.$$

Since $\mathbb{F}$ is a field, for every Boolean assignment a,

$$\llbracket D \rrbracket(a) = 0 \quad \iff \quad f_i(a) = 0 \text{ for some } i \in [k] \quad \iff \quad a \models D.$$

Thus, $\llbracket D \rrbracket = 0$ represents the same set of Boolean assignments as the line D.

Each polynomial f_i is represented by a $\Sigma\Pi$ circuit, so $\llbracket D \rrbracket$ has a $\Sigma\Pi\Sigma\Pi$ representation with a single top product term. Across the entire derivation, we reuse all shared polynomial subcircuits and introduce at most a product gate and a unary sum gate for each distinct proof line. Thus the translated proof has total circuit size $O(\text{size}(\pi))$.

If D is a semantic consequence of earlier lines $D_1, \ldots, D_t$, then the equivalence above shows that $\llbracket D \rrbracket = 0$ is a semantic consequence of

$$\llbracket D_1 \rrbracket = 0, \ldots, \llbracket D_t \rrbracket = 0.$$

Thus, every inference translates into a valid semantic $\Sigma\Pi\Sigma\Pi$-PC inference of the same arity. Initial clauses translate into their standard polynomial encodings, and the empty disjunction translates into the empty product 1. This gives the required refutation π'.

Every translated line contains a single top product term, and hence $S(\pi') \leq L$. Finally, if every f_i has degree at most w, then every bottom product gate in its $\Sigma\Pi$ representation has fan-in at most w. Therefore, the simulation maps Res(PC_w) into $\Sigma\Pi\Sigma\Pi^{[w]}$-PC. □

3 Warm-Up: ΣΠΣ-PC Lower Bound for the Bit Pigeonhole Principle

In this section, we prove that semantic ΣΠΣ-PC requires exponential size to refute the *bit pigeonhole principle* BPHP, which we recall first. The proof in this section is best viewed as a simplified presentation of the Res(⊕) lower bound discovered by Braun [16].

Definition 3.1 (bit pigeonhole principle BPHP)**.** Let $n = 2^\ell$ and $m > n$, and introduce variables $y_{u,t}$ for pigeons $u \in [m]$ and bit positions $t \in [\ell]$. For every pair of pigeons $u \neq u' \in [m]$ and every hole address $a \in \{0,1\}^\ell$, the *collision clause*

$$D_{u,u',a} = \bigvee_{t=1}^{\ell} (y_{u,t} \neq a_t) \vee \bigvee_{t=1}^{\ell} (y_{u',t} \neq a_t) \tag{1}$$

asserts that u and u' cannot both be mapped to the hole a. The formula BPHP_n^m is the conjunction

$$\mathsf{BPHP}_n^m := \bigwedge_{u \neq u' \in [m]} \bigwedge_{a \in \{0,1\}^\ell} D_{u,u',a},$$

asserting that no two pigeons collide in any hole. Note that no pigeon axioms are needed because every bit string already names a hole. Equivalently, putting

$$\delta_a(y_u) = \prod_{t=1}^{\ell} (a_t y_{u,t} + (1 - a_t)(1 - y_{u,t})), \text{ where } u \in [m] \text{ and } a \in \{0,1\}^\ell,$$

the polynomial encoding of BPHP_n^m consists of the equations

$$\delta_a(y_u)\delta_a(y_{u'}) = 0 \qquad \forall u \neq u' \in [m], a \in \{0,1\}^\ell.$$

Theorem 3.2. *Let $n = 2^\ell$ for a sufficiently large integer ℓ, and let $m > n$. Any semantic ΣΠΣ-PC refutation of BPHP_n^m over $\mathbb{F}_2$ requires size $2^{\Omega(n^2/(m\ell^2))}$.*

By the simulation results we have seen in the previous section, this immediately implies Braun's result that any Res(⊕) refutation of BPHP_n^m requires exponential size.

First, we will introduce some notation. Henceforth, all polynomials are over $\mathbb{F}_2$, and we write $\mathcal{B}_n$ for $\mathcal{B}_n(\mathbb{F}_2)$. Let $\mathcal{B}_n^{\leq k}$ denote the subspace of multilinear polynomials of degree at most k. Importantly, we will be interpreting all polynomials $p \in \mathbb{F}_2[x_1, \ldots, x_n]$ as elements of $\mathcal{B}_n$, and hence $\deg(p)$ will represent the degree of the *multilinear representation* of p. Recall that $p \equiv q$ if $p = q \bmod \langle x_i^2 - x_i \rangle_{i=1}^n$, and note that $p \equiv q$ if and only if $p(x) = q(x)$ for all $x \in \{0,1\}^n$. The lines of a semantic ΣΠΣ proof are ΣΠΣ circuits, and the primary complexity measure we will use for these lines is recorded in the next definition.

Definition 3.3. Suppose $f = \prod_{j=1}^s L_j$ is a ΠΣ circuit over $\mathbb{F}_2$, noting that each L_j is an affine form. Define the *rank* of f to be

$$\operatorname{rank}(f) := \begin{cases} 0 & f \equiv 0 \\ \dim(\operatorname{span}\{L_j + 1 : j \in [s]\}) & \text{otherwise} \end{cases}$$

If $g = \sum_i f_i$ is a ΣΠΣ circuit over $\mathbb{F}_2$, then we define $\operatorname{rank}(g) = \max_i \operatorname{rank}(f_i)$. If π is a ΣΠΣ-PC proof, then $\operatorname{rank}(\pi)$ is the maximal rank of any line in the proof.

For intuition, over $\mathbb{F}_2$, any ΠΣ circuit f encodes an indicator function for a (possibly empty) affine space defined by its underlying affine forms. This means that if $f = \prod_{j=1}^s L_j$ is not identically 0 as a Boolean function, then $\operatorname{rank}(f)$ is exactly the codimension of the affine space defined by $L_1 = L_2 = \cdots = L_s = 1$. For later, an important observation is that if f is a ΣΠΣ circuit, then $\deg(f) \leq \operatorname{rank}(f)$, where $\deg(f)$ denotes the degree of the multilinear polynomial representing f.

The high-level idea of the proof of Theorem 3.2 is analogous to the random restriction arguments used for Resolution lower bounds. In simplified form, the argument proceeds in two steps:

1. First, we design a low-degree polynomial p which, when multiplied by all lines of the $\Sigma\Pi\Sigma$-PC refutation π, will yield a simplified refutation π' with all top product terms of high rank removed (i.e. set to 0), while preserving the axioms of BPHP.

2. Second, since all remaining product terms have low rank, we can simulate π' by a semantic PC refutation of low degree. From this, we can appeal to known degree lower bounds for BPHP to obtain a final contradiction.

Usually in proof complexity a "random restriction" is a partial restriction ρ to the variables of the underlying formula chosen from some distribution, and it is not hard to see that $\Pi\Sigma$ circuits are resistant to being simplified by such a simple restriction. In this argument, we choose a more complicated object — a low-degree polynomial — by which we "restrict" (i.e. multiply) the lines of our proof by in order to eliminate all the top product terms of high rank. We isolate this as a definition:

Definition 3.4. If $\mathcal{S} \subseteq \mathcal{B}_n$ and $p \in \mathcal{B}_n$, then p *annihilates*[4] $\mathcal{S}$ if $fp \equiv 0$ for all $f \in \mathcal{S}$.

Multiplying through the entire refutation by this polynomial p will remove the high-rank top product terms, but then we are left with a low-rank *derivation* of the polynomial p from the axioms of BPHP_n^m, which we must rule out. Thankfully, as we noted above, if f is a $\Sigma\Pi\Sigma$ circuit, then $\deg(f) \leq \operatorname{rank}(f)$, and hence the proof π' is actually a low-degree semantic polynomial calculus proof. From here, we can apply a second restriction argument, and use known polynomial calculus degree lower bounds for BPHP_n^m to obtain a final contradiction. To make this second restriction argument work, we must ensure that the annihilating polynomial has a special form:

Definition 3.5. A multilinear polynomial p over the variables of BPHP_n^m is *pigeon-linear* if every monomial of p mentions, for each pigeon $u \in [m]$, at most one variable of the form $y_{u,t}$. Equivalently, p is pigeon-linear if

$$p \in H := \operatorname{span}\left\{ \prod_{i \in I} y_{i,t_i} : I \subseteq [m], \forall i \in I : t_i \in [\ell] \right\}.$$

We let $H^{\leq k} \subseteq H$ denote the subspace of pigeon-linear polynomials of degree at most k.

The analogue of the restriction argument is the next lemma, which is essentially [16, Lemma 6.4].

Lemma 3.6. *Let $1 \leq k \leq m$ be a positive integer parameter, and recall $n = 2^\ell$. Suppose $\mathcal{S}$ is a set of $\Pi\Sigma$ circuits with $|\mathcal{S}| \leq S$ over the variables of BPHP_n^m such that for all $f \in \mathcal{S}$, $\operatorname{rank}(f) > 3\ell(k+1)$. If $S \leq \exp(k^2/m)/4$, then there is a non-zero pigeon-linear polynomial $p \in H^{\leq k}$ annihilating $\mathcal{S}$.*

Proof. If $3(k+1) \geq m$, then every $\Pi\Sigma$ circuit has rank at most $m\ell \leq 3\ell(k+1)$, so $\mathcal{S} = \varnothing$ and we may take $p = 1$. We may therefore assume that $3(k+1) < m$ and proceed by a counting argument. Note that, as a vector space,

$$\dim H^{\leq k} = \sum_{i=0}^{k} \binom{m}{i} \ell^i.$$

This is because each basis monomial is obtained by selecting a set of i pigeons, and choosing one of ℓ bits associated with each of the i selected pigeons to appear in the monomial. So, define $N = \dim H^{\leq k}$, and note that every polynomial $p \in H^{\leq k}$ is specified by assigning $\{0,1\}$-values to the N coefficients of these monomials.

Assume without loss of generality that $\mathcal{S}$ does not contain any polynomials $p \equiv 0$. Let $f \in \mathcal{S}$ be a $\Pi\Sigma$ circuit with $\operatorname{rank}(f) = r > 3\ell(k+1)$. Note that f is an indicator of an affine subspace of $\{0,1\}^{m\ell}$ of codimension $r > 3\ell(k+1)$, and hence we can parametrise this subspace by a collection of $m\ell - r$ variables $z_1, z_2, \ldots, z_{m\ell-r}$. Formally, this means that there is a vector a and a matrix A such that

$$f^{-1}(1) = \{x \in \{0,1\}^{m\ell} : \exists z \in \{0,1\}^{m\ell-r} : x = a + Az\}.$$

[4]The term "annihilation" comes from ring theory: if R is a commutative ring and $S \subseteq R$, then the *annihilator* of S, denoted $\operatorname{Ann}(S)$, is the set of elements $a \in R$ such that $as = 0$ for all $s \in S$.

Therefore, given any polynomial $p \in H^{\leq k}$, the condition $fp \equiv 0$ is implied by the stronger condition that the polynomial $p(a + Az) \equiv 0$. Since $p \in H^{\leq k}$, $\deg p \leq k$, and, importantly, an affine substitution cannot increase degree, so $\deg p(a + Az) \leq k$.

Now we may perform the counting argument. Observe that there are at most $\binom{m\ell-r+k}{k}$ monomials over variables $z_1, \ldots, z_{m\ell-r}$ with degree at most k. Hence, in order for $p(a + Az) \equiv 0$, we must enforce at most $\binom{m\ell-r+k}{k}$ homogeneous linear equations on the N coefficients of p corresponding to the polynomial f. Setting $r = 3\ell(k + 1) + 1$ and $d = m\ell - r$, then

$$\binom{d+k}{k} \leq \frac{(d+k)^k}{k!}, \quad N \geq \binom{m}{k}\ell^k \geq \frac{(\ell(m-k+1))^k}{k!}.$$

Hence

$$\frac{\binom{d+k}{k}}{N} \leq \left(\frac{d+k}{\ell(m-k+1)}\right)^k.$$

Substituting the value of r into d yields $d + k \leq \ell(m - 2k)$. Continuing the previous calculation, we get

$$\begin{aligned}\left(\frac{d+k}{\ell(m-k+1)}\right)^k &\leq \left(\frac{\ell(m-2k)}{\ell(m-k+1)}\right)^k \\ &\leq \left(\frac{m-2k}{m-k+1}\right)^k \\ &\leq \left(1 - \frac{k}{m}\right)^k \leq \exp(-k^2/m).\end{aligned}$$

Therefore, each $f \in \mathcal{S}$ contributes at most $N \exp(-k^2/m)$ linear constraints on the coefficients of p. So, by subadditivity of rank, enforcing $fp \equiv 0$ for all $f \in \mathcal{S}$ requires adding at most

$$NS \exp(-k^2/m) \leq \frac{N}{4}$$

linearly independent equations to the coefficients of p. Therefore, we can choose a polynomial $p \in H^{\leq k}$ which is non-zero and satisfies all of these equations, since $\dim H^{\leq k} = N$ and the dimension of the restricted space is at least $N - N/4 = 3N/4$. □

With this lemma we can complete the proof. Recalling that $S(\pi)$ is the *top-product size* of π, if $S(\pi)$ is small, then multiplying every line of π by p will result in a low-degree PC derivation of p from the axioms. Afterwards, we use a further restriction argument to derive a contradiction.

We will need the following polynomial calculus degree lower bound for BPHP^m_n, which is easily proved by reduction to Razborov's degree lower bound for the functional weak pigeonhole principle in the unary encoding [59] (cf. Appendix A).

Theorem 3.7. *For all sufficiently large $n = 2^\ell$ and any $m > n$, any semantic polynomial calculus refutation of BPHP^m_n over $\mathbb{F}_2$ requires degree $n/6$.*

We can now prove Theorem 3.2.

Proof of Theorem 3.2. By Theorem 3.7, for all sufficiently large ℓ with $n = 2^\ell$ and any $m > n$, any semantic polynomial calculus refutation of BPHP^m_n requires degree at least $n/6$. Towards a contradiction, set $k = \lfloor n/(192\ell) \rfloor$, let $m > n$, suppose there is a semantic $\Sigma\Pi\Sigma$-PC refutation π of BPHP^m_n with $S(\pi) \leq \exp(k^2/m)/4$, and let $\mathcal{S}$ denote the set of top-product polynomials in π. By Lemma 3.6, there is a non-zero, pigeon-linear polynomial p with $\deg p = t \leq k$ such that p annihilates all polynomials in $\mathcal{S}$ with rank $> 3\ell(k + 1)$. Let π' be obtained from π by multiplying every line by p, noting that this can be obtained by starting from the axiom lines of π, multiplying them by p, and then carrying out the operations of the proof in the usual way. Since every line in π is a sum of top product terms, multiplying it by p removes every term of rank greater than $3\ell(k + 1)$. Setting $d = 3\ell(k + 1)$, all surviving polynomials in the proof therefore have degree at most $R = d + k \leq 4\ell(k + 1)$. Hence, the resulting proof π' is a semantic PC proof of the polynomial p in degree $R = O(\ell k)$.

We now prove this is impossible, by reduction to Theorem 3.7. We first give the idea of the proof. Choose any maximum-degree monomial $M = \prod_{j=1}^{t} y_{i_j,a_j}$ in p, and write $p = M + p'$. Since p is pigeon-linear, all queried pigeons in M are distinct. The idea is to consider, for each $\alpha \in \{0,1\}^t$, the restriction ρ_α which assigns $y_{i_j,a_j} = \alpha_j$ and leaves all other variables unassigned. Consider the following new proof: for every α, we hit π' with ρ_α, obtaining $\pi' \restriction \rho_\alpha$. This will be a proof of $p \restriction \rho_\alpha$ from restricted axioms of $\mathsf{BPHP}^m_n \restriction \rho_\alpha$. Summing the final lines of all of these proofs up, we will obtain a proof of the polynomial $1 + \sum_\alpha p' \restriction \rho_\alpha$. By some careful massaging of the restrictions ρ_α, we can ensure the contribution from p' is 0, and the restricted axioms used come from copies of a slightly smaller copy of the pigeonhole principle. Hence, we will have a low-degree polynomial calculus refutation of a smaller copy of BPHP, which will violate the existing degree lower bounds in Theorem 3.7.

Let us now proceed more formally. Consider any maximum-degree monomial $M = \prod_{j=1}^{t} y_{i_j,a_j}$ in p. Let

$$H = \{z \in \{0,1\}^\ell : z_1 = z_2 = 0\}.$$

For each pigeon i_j appearing as an index in M, select two hole labels $\alpha^{j,0}, \alpha^{j,1}$ such that for each j:

- $\alpha^{j,0}, \alpha^{j,1}$ differ only on bit a_j, with $\alpha^{j,b}_{a_j} = b$ for $b \in \{0,1\}$,
- both $\alpha^{j,0}, \alpha^{j,1} \notin H$, and
- the hole labels chosen for each pigeon are unique among all pigeons mentioned by M.

It is straightforward to show such a choice exists. Consider the pigeons $i_1, i_2, \ldots, i_t$ in order. For each $j \in [t]$, if $a_j \in \{1,2\}$, then set the coordinate in $\{1,2\} \setminus \{a_j\}$ to 1 in $\alpha^{j,0}, \alpha^{j,1}$, and choose the first available label for the remaining coordinates, leaving a_j free. Otherwise, if $a_j \notin \{1,2\}$, then set $\alpha^{j,b}_1 = 1$, and choose the first available label for each coordinate other than a_j. For each j, this gives $n/4$ disjoint candidate pairs of hole labels outside H. The $2(j-1)$ previously used labels exclude at most $2(j-1)$ candidate pairs, so an unused pair exists since $j \le t \le k \le n/8$.

Now, for each $\beta \in \{0,1\}^t$, consider the following partial restriction ρ_β. For each pigeon i_j queried by M, ρ_β assigns i_j to the hole α^{j,β_j}. For each pigeon i not queried by M, ρ_β assigns the first two bits to 0, and leaves all other bits unassigned. Under ρ_β, every collision axiom involving a queried pigeon is satisfied, and the remaining pigeons are restricted to the $n/4$ holes in H. Thus the restricted formula is $\mathsf{BPHP}^{m-t}_{n/4}$, after deleting satisfied axioms and identifying the remaining $\ell - 2$ bits with hole addresses. Moreover, $m - t \ge m - k > n/4$. Moreover, restrictions can only decrease the degree, and hence we have for each β a proof of $p \restriction \rho_\beta$ from $\mathsf{BPHP}^{m-t}_{n/4}$ in degree at most R.

Consider the proof obtained by combining all these proofs together, deriving $p \restriction \rho_\beta$ for each β, and then deriving $\sum_\beta p \restriction \rho_\beta$. We claim $\sum_\beta p \restriction \rho_\beta = 1$. To see this, observe that $M \restriction \rho_\beta = 0$ unless $\beta = 1^t$, in which case it evaluates to 1. Let M' be any other monomial in p. If M' omits at least one pigeon i_j chosen by M, then we get an even number of copies of M' by pairing up $\alpha^{j,0}, \alpha^{j,1}$ as β ranges over $\{0,1\}^t$. Otherwise, suppose M' depends on all pigeons chosen by M. In this case, since $M' \ne M$, there must be a pigeon i_j such that M' contains the variable $y_{i_j,b}$ for some $b \ne a_j$, and by pigeon linearity this is the *only* variable from pigeon i_j in M'. Hence we may again pair up $\alpha^{j,0}, \alpha^{j,1}$ as we range over $\beta \in \{0,1\}^t$, as the bit $y_{i_j,b}$ will not depend on flipping the bit a_j, cancelling out the monomial M'.

Thus we are left with a degree at most $R \le 4\ell(k+1)$ refutation of $\mathsf{BPHP}^{m-t}_{n/4}$. Applying Theorem 3.7 with $n' = n/4$ gives $R \ge n/24$. However, our choice of k gives

$$R \le 4\ell(k+1) \le \frac{n}{48} + 4\ell < \frac{n}{24}$$

for sufficiently large ℓ, a contradiction.

Therefore, every semantic $\Sigma\Pi\Sigma$-PC refutation π of BPHP^m_n satisfies $S(\pi) > \exp(k^2/m)/4$. For sufficiently large ℓ, we have $k \ge n/(384\ell)$, and hence

$$S(\pi) \ge \frac{1}{4} \exp\left(\frac{n^2}{384^2 m \ell^2}\right).$$

Together with the trivial bound $S(\pi) \ge 2$, this gives $S(\pi) = 2^{\Omega(n^2/(m\ell^2))}$. □

4 Abstract Lower Bound Framework: Degree Reduction

Here we generalise the argument of Section 3. We combine it with approximation techniques for $\mathsf{AC}^0[2]$ circuits of Razborov and Smolensky [60, 64]. This gives an abstract technique for proving size lower bounds for semantic refutations whose lines are $\Sigma\Pi\Sigma$ or $\Sigma\Pi\Sigma\Pi$ circuits.

Notation. Fix a set of polynomial axioms $\mathcal{F} \subseteq \mathcal{B}_n$ and let $\mathcal{F}_0 \subseteq \mathcal{F}$ be a subset that we shall treat as "background axioms" (this will be helpful notation later when we are considering constructions, but for the purpose of the overview, we can regard $\mathcal{F}_0 = \emptyset$). Define

$$\mathcal{B}_{\mathcal{F}_0} = \mathcal{B}_n/\langle \mathcal{F}_0 \rangle,$$

so that $\mathcal{B}_\varnothing = \mathcal{B}_n$. Recall that the Boolean equations already hold in $\mathcal{B}_n$, and so passing to $\mathcal{B}_{\mathcal{F}_0}$ imposes the additional equations in $\mathcal{F}_0$. By the known equivalence between $\mathcal{B}_n$ and Boolean functions over $\mathbb{F}_2$, we can further identify $\mathcal{B}_{\mathcal{F}_0}$ with the algebra of Boolean functions on the set

$$\Omega = \{x \in \{0,1\}^n : g(x) = 0 \text{ for every } g \in \mathcal{F}_0\}.$$

For $q \in \mathbb{F}_2[x_1, \ldots, x_n]$, let $\bar{q}$ denote its image in $\mathcal{B}_{\mathcal{F}_0}$. This image is obtained by first taking the quotient by the Boolean equations and then by $\langle \mathcal{F}_0 \rangle$ in $\mathcal{B}_n$. For any $p \in \mathcal{B}_{\mathcal{F}_0}$, define

$$\deg p := \min\{\deg q : q \in \mathbb{F}_2[x_1, \ldots, x_n], \bar{q} = p\}.$$

In other words, the degree of $p \in \mathcal{B}_{\mathcal{F}_0}$ is the minimum total degree among all polynomials in $\mathbb{F}_2[x_1, \ldots, x_n]$ whose residue class is that element. The original polynomial remains a representative after imposing the background equations. Since degree in the quotient $\mathcal{B}_{\mathcal{F}_0}$ is the minimum degree of a representative, passing to the quotient cannot increase degree.

We denote by $\mathcal{B}_{\mathcal{F}_0}^{\leq k}$ the subspace of $\mathcal{B}_{\mathcal{F}_0}$ consisting of polynomials of degree at most k, and we define

$$H_{\mathcal{F}_0}(k) := \dim(\mathcal{B}_{\mathcal{F}_0}^{\leq k})$$

to count the number of polynomials in a basis for $\mathcal{B}_{\mathcal{F}_0}^{\leq k}$.[5]

When $\mathcal{F}_0 = \varnothing$, we also write

$$H_n(k) = \dim(\mathcal{B}_n^{\leq k}) = \sum_{j=0}^{\min\{k,n\}} \binom{n}{j}.$$

From now on, we regard each polynomial as an element of $\mathcal{B}_{\mathcal{F}_0}$ and keep the same notation. Thus two polynomials are treated as equal whenever their difference belongs to the ideal $\langle \mathcal{F}_0 \rangle$. We also keep the notation $\mathcal{F}$ for the input set of axioms. Unless stated otherwise, PC derivations take place in $\mathcal{B}_{\mathcal{F}_0}$, and semantic inferences are interpreted on Ω. The axioms in $\mathcal{F}_0$ are zero in this quotient, so all their ideal consequences are available without a derivation. Every semantic refutation remains valid modulo $\mathcal{F}_0$, and its number of distinct non-zero top product terms cannot increase (because keeping the original circuit representation of every line gives a semantic refutation modulo the background axioms. Some top product terms may become zero, and distinct terms may become equal).

Reduction and pseudo-reduction operators. Let $\mathcal{F} = \{f_1, \ldots, f_s\} \subseteq \mathcal{B}_{\mathcal{F}_0}^{\leq d_0}$ be the set of the initial polynomials, considered in the quotient $\mathcal{B}_{\mathcal{F}_0}$, and fix a degree parameter k, which we assume to be at most n and at least d_0, an upper bound on the degrees of the polynomials in $\mathcal{F}$.

If $I = \langle \mathcal{F} \rangle$ is the ideal generated by the polynomials in $\mathcal{F}$, then one can think of the natural *reduction operator* modulo this ideal: every element $p \in \mathcal{B}_{\mathcal{F}_0}$ is reduced to a remainder in $\mathcal{B}_{\mathcal{F}_0}$ that represents its residue class in $\mathcal{B}_{\mathcal{F}_0}/I$. In other words, reducing p means considering p modulo I, and then choosing a normal-form representative. In particular, a reduction operator sends every element of

[5] $H_{\mathcal{F}_0}$ is the affine Hilbert function of the ideal generated by $\mathcal{F}_0$ and the Boolean axioms.

I to 0. Since the families of polynomial equations we consider are unsatisfiable, the ideal $\langle \mathcal{F} \rangle$ contains 1, and hence contains *all* polynomials, and so reducing modulo this ideal will just send everything to 0.

For the sake of proving lower bounds, we therefore use an operator which is "like" reduction modulo this ideal, but only up to polynomials of degree at most k, for some parameter k. Clegg, Edmonds, and Impagliazzo [23] already characterise bounded-degree PC derivations by a linear space and give a corresponding reduction procedure. Related linear maps occur in the earlier design method for Nullstellensatz degree lower bounds [10]; see Buss [17] for a general formulation. Razborov [59] developed the reduction operator method for PC degree lower bounds. We define such an operator formally:

Definition 4.1 (Pseudo-reduction operator)**.** Let $\mathcal{F}_0$ be any set of background axioms, and let W be any linear space. For any family of polynomials $\mathcal{F}$, a linear map

$$\mathcal{R} \colon \mathcal{B}_{\mathcal{F}_0}^{\leq k} \to W$$

is a *degree-k pseudo-reduction operator* for $\mathcal{F}$ if

$$\mathcal{R}(f) = 0 \qquad \text{for every } f \in \mathcal{F},$$

and, for every $p \in \mathcal{B}_{\mathcal{F}_0}^{\leq k}$ and every variable x_i such that $\deg(x_i p) \leq k$,

$$\mathcal{R}(p) = 0 \quad \Longrightarrow \quad \mathcal{R}(x_i p) = 0. \tag{2}$$

It is helpful to think of the kernel of a pseudo-reduction operator as a kind of "pseudo-ideal": it is a linear subspace of $\mathcal{B}_{\mathcal{F}_0}^{\leq k}$ containing $\mathcal{F}$ and closed under multiplication by a variable whenever the resulting degree is at most k. In particular, this kernel contains every polynomial with a PC derivation of degree at most k, but it may contain other polynomials as well. Unlike a true reduction operator, this operator need not compute normal-form remainders.

As one would expect, the existence of pseudo-reduction operators immediately implies PC degree lower bounds. Let

$$D_k(\mathcal{F}) \subseteq \mathcal{B}_{\mathcal{F}_0}^{\leq k} \tag{3}$$

denote the set of all polynomials that have a (syntactic) PC derivation from $\mathcal{F}$, in which every polynomial has degree at most k. In the next lemma, $\operatorname{rank}(\mathcal{R}) = \dim \operatorname{im}(\mathcal{R})$ is the dimension of its image. We have the following simple lemma.

Lemma 4.2 (Pseudo-reduction lemma)**.** *Let $\mathcal{R}$ be a degree-k pseudo-reduction operator for $\mathcal{F}$. Then:*

1. *Every polynomial $p \in D_k(\mathcal{F})$ satisfies $\mathcal{R}(p) = 0$.*
2. *If $\operatorname{rank}(\mathcal{R}) > 0$ (i.e., $\operatorname{im}(\mathcal{R}) \neq \{0\}$), then $\mathcal{F}$ has no PC refutations of degree at most k.*

Proof Sketch. The first point follows by induction over the proof of p using Equation (2). To see the second point, assume $\operatorname{rank}(\mathcal{R}) > 0$, and by contradiction suppose that $\mathcal{F}$ has a PC refutation of degree $\leq k$. This implies that $1 \in D_k(\mathcal{F})$, and so $\mathcal{R}(1) = 0$. But then Equation (2) and linearity imply that *all* degree-k polynomials p satisfy $\mathcal{R}(p) = 0$, and so $\dim \mathcal{B}_{\mathcal{F}_0}^{\leq k} = \dim \ker \mathcal{R}$, which implies $\operatorname{rank}(\mathcal{R}) = 0$, a contradiction. □

The lemma is "complete" in the following sense: we can always find a degree-k pseudo-reduction operator whose kernel is *exactly* the set of polynomials derivable in degree-k polynomial calculus. We will use this construction for the functional pigeonhole principle in Theorem 6.12.

Proposition 4.3. *For any $\mathcal{F}$ there is a degree-k pseudo-reduction operator $\mathcal{R}$ with $\ker(\mathcal{R}) = D_k(\mathcal{F})$.*

Proof. It is easy to see that $D_k(\mathcal{F})$ is a linear subspace of $\mathcal{B}_{\mathcal{F}_0}^{\leq k}$, since the polynomials derivable in degree $\leq k$ polynomial calculus are closed under linear combination. It contains $\mathcal{F}$ and is closed under every multiplication step that stays within the degree bound:

$$p \in D_k(\mathcal{F}) \text{ and } \deg(x_i p) \leq k \quad \Longrightarrow \quad x_i p \in D_k(\mathcal{F}).$$

This means we can quotient by it — define $\mathcal{R} : \mathcal{B}_{\mathcal{F}_0}^{\leq k} \to \mathcal{B}_{\mathcal{F}_0}^{\leq k} / D_k(\mathcal{F})$ by the quotient map, so $\mathcal{R}(p)$ is exactly the residue class of p modulo $D_k(\mathcal{F})$. □

When $\mathcal{F}_0 = \varnothing$, the method of Alekhnovich and Razborov [2] provides explicit operators of this kind, usually with $W = \mathcal{B}_n^{\leq k}$, which simplify polynomials without increasing their degree. Their defining identities mirror the PC inference rules and, in particular, imply Equation (2).[6] Filmus [29] gave a complementary abstract account of the Alekhnovich–Razborov method. Mikša and Nordström [54] developed a general framework for such operators, which we use below. Their framework likewise sets aside a collection $\mathcal{F}_0$ of axioms that are assumed throughout the reduction, so our definition of $\mathcal{B}_{\mathcal{F}_0}$ is compatible with theirs.

Establishing size-lower bounds using pseudo-reduction operators. Lemma 4.2 shows that pseudo-reduction operators imply degree lower bounds for PC, as long as they have non-zero rank. Our main theorems, stated next, generalise this simple observation, and provide a *quantitative* relationship between the rank of the operator and size lower bounds in ΣΠΣ-PC and $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutations.

Our first theorem holds for ΣΠΣ-PC, and as none of our applications of it require fixing a set of background axioms, we do not state it as such. To state our theorem, recall that if π is a ΣΠΣ-PC proof, then $S(\pi)$ counts the number of distinct polynomials appearing as top-products in the circuits occurring in π. This theorem is proved by directly generalising the arguments from Section 3:

Theorem 4.4 (Size bound for ΣΠΣ-PC via a reduction operator)**.** *Let $\mathcal{F} \subseteq \mathcal{B}_n$, and let d_0 denote the maximum degree of the polynomials in $\mathcal{F}$. Fix $1 \leq k, r \leq n$, and put $B = 3\big(k + \max\{d_0, r\}\big)$. Suppose that $B \leq n$ and that $\mathcal{R}\colon \mathcal{B}_n^{\leq B} \to W$ is a degree-B pseudo-reduction operator for $\mathcal{F}$, and write $\mathcal{R}_{\leq k}$ for the restriction of $\mathcal{R}$ to $\mathcal{B}_n^{\leq k}$. Then every semantic ΣΠΣ-PC refutation π of $\mathcal{F}$ satisfies*

$$S(\pi) \geq \frac{\operatorname{rank}(\mathcal{R}_{\leq k})}{2\binom{n-r}{k}}.$$

Our next theorem provides lower bounds for $\Sigma\Pi\Sigma\Pi$-PC in terms of $\operatorname{rank}(\mathcal{R})$, and for this we need to use the additional structure of the background axioms $\mathcal{F}_0$. This will require us to slightly modify the definition of $S(\pi)$. Consider a proof π in $\Sigma\Pi\Sigma\Pi$-PC. Fix the circuits representing the proof lines of π, and consider the residue classes in $\mathcal{B}_{\mathcal{F}_0}$ of the polynomials computed by their top product gates. We denote by $S(\pi)$ the number of distinct non-zero classes among them. (This count uses the given circuit representations; it does not minimise over alternative representations.)

The background axioms $\mathcal{F}_0$ are applied using the notion of *n-divisions*, introduced by Impagliazzo, Pudlák, and Sgall [40]. Let $N = nh$ be the total number of variables, partitioned into disjoint sets $V_1, \ldots, V_n$, each of size h. For every $i \in [n]$, introduce the set $E_i = \{xy : x, y \in V_i,\ x \neq y\}$ and let $\mathcal{F}_0 = \bigcup_{i\in[n]} E_i$.[7] The polynomial system $\mathcal{F} \subseteq \mathcal{B}_N$ has an *n-division* if $\mathcal{F}_0 \subseteq \mathcal{F}$.

Theorem 4.5 (Rank bound for $\Sigma\Pi\Sigma\Pi^{[w]}$-PC with n-divisions)**.** *Let $\mathcal{F} \subseteq \mathcal{B}_N$ have an n-division $V_1, \ldots, V_n$, where every part has size h and $N = nh$. Let $\mathcal{F}_0 \subseteq \mathcal{F}$ be the set of all products xy of distinct variables from the same part. Consider $\mathcal{F}$ in the quotient $\mathcal{B}_{\mathcal{F}_0} = \mathcal{B}_N/\langle\mathcal{F}_0\rangle$, and suppose that its axioms have degree at most d_0 there. Fix integers $w \geq 1$, $2w \leq k \leq n$, and $r \geq 0$, and put*

$$B = 3\big(k + \max\{d_0, rw\}\big).$$

Suppose that $B \leq n$ and that $\mathcal{R}\colon \mathcal{B}_{\mathcal{F}_0}^{\leq B} \to W$ is a degree-B pseudo-reduction operator for $\mathcal{F}$. Write $\mathcal{R}_{\leq k}$ for the restriction of $\mathcal{R}$ to $\mathcal{B}_{\mathcal{F}_0}^{\leq k}$. Then every semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutation π of $\mathcal{F}$ satisfies

$$S(\pi) \geq \frac{\operatorname{rank}(\mathcal{R}_{\leq k})}{2\binom{n}{k}h^k} \exp\left(\frac{r}{4}\left(\frac{k}{2N}\right)^w\right).$$

Every polynomial system on n variables has an n-division into singleton parts, with $h = 1$ and $N = n$. There are then no pairs of distinct variables within a part, so $\mathcal{F}_0 = \varnothing$. Applying Theorem 4.5 with this division gives the following version without background axioms.

[6] Pseudo-reduction operators are often required to satisfy the stronger identity $\mathcal{R}(x_i p) = \mathcal{R}\big(x_i \mathcal{R}(p)\big)$, which implies (2).

[7] It is helpful to think of this $\mathcal{F}_0$ as the hole axioms of the pigeonhole principle: the sets V_i represent n holes, and for each hole there are exactly h pigeons allowed to fly into this hole.

Corollary 4.6 (Rank bound for $\Sigma\Pi\Sigma\Pi^{[w]}$-PC without assuming n-divisions)**.** *Let $\mathcal{F} \subseteq \mathcal{B}_n$ be a polynomial system whose axioms have degree at most d_0. Fix integers $w \geq 1$, $2w \leq k \leq n$, and $r \geq 0$, and put*

$$B = 3(k + \max\{d_0, rw\}).$$

Suppose that $B \leq n$ and that $\mathcal{R} \colon \mathcal{B}_n^{\leq B} \to W$ is a degree-B pseudo-reduction operator for $\mathcal{F}$. Write $\mathcal{R}_{\leq k}$ for the restriction of $\mathcal{R}$ to $\mathcal{B}_n^{\leq k}$. Then every semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutation π of $\mathcal{F}$ satisfies

$$S(\pi) \geq \frac{\operatorname{rank}(\mathcal{R}_{\leq k})}{2\binom{n}{k}} \exp\left(\frac{r}{4}\left(\frac{k}{2n}\right)^w\right).$$

4.1 Proof of the $\Sigma\Pi\Sigma$-PC size bound

We will begin by proving Theorem 4.4, as it is simpler, and is essentially just a generalisation of the counting technique from Section 3. Let $\mathcal{F} \subseteq \mathcal{B}_n$ with maximum degree at most d_0, and fix $1 \leq k, r \leq n$ with $B = 3(k + \max\{d_0, r\}) \leq n$. Suppose that $\mathcal{R} : \mathcal{B}_n^{\leq B} \to W$ is a degree-B pseudo-reduction operator for $\mathcal{F}$, and write $\mathcal{R}_{\leq k}$ for the restriction of $\mathcal{R}$ to $\mathcal{B}_n^{\leq k}$.

Recall from Section 3 that if $f = \Pi_{j=1}^s L_j$ is a $\Pi\Sigma$ circuit, then

$$\operatorname{rank}(f) := \begin{cases} 0 & f \equiv 0 \\ \dim(\operatorname{span}\{L_j + 1 : j \in [s]\}) & \text{otherwise} \end{cases}$$

If π is a $\Sigma\Pi\Sigma$-PC proof, then $\operatorname{rank}(\pi)$ is the maximal rank of any line in the proof. If f is not identically 0 as a Boolean function, then $\operatorname{rank}(f)$ is exactly the codimension of the affine space defined by $L_1 = L_2 = \cdots = L_s = 1$, which is a crucial complexity measure of the lines. Choosing independent equations defining this affine space expresses its indicator as a product of $\operatorname{rank}(f)$ affine forms, so $\deg(f) \leq \operatorname{rank}(f)$.

Overall, the structure of the proof of Theorem 4.4 will be the same. We will use a counting argument, analogous to Lemma 3.6, to construct a low-degree polynomial f which annihilates all terms of large codimension. The main place where our argument differs is that we will construct a low-degree annihilator f which satisfies $\mathcal{R}(f) \neq 0$. Then, after multiplying through by f, all remaining terms will have low degree. This means we can transform the proof into a low-degree PC_2 derivation of f from $\mathcal{F}$, but f *cannot* have such a derivation since $\mathcal{R}(f) \neq 0$. This completes the proof.

We now begin in earnest. First, we generalise Lemma 3.6. The assumption $3(k + \max\{d_0, r\}) \leq n$ implies $k \leq (n - r)/3$, so we use this simpler sufficient condition instead. In our application of the lemma, we will take the linear map $L = \mathcal{R}_{\leq k}$, while Lemma 3.6 can be viewed as using the identity map for L.

Lemma 4.7. *Let $1 \leq k, r \leq n$, and let $L \colon \mathcal{B}_n^{\leq k} \to W$ be a linear map. Suppose that $\mathcal{S}$ is a finite set of $\Pi\Sigma$ circuits on n variables such that $\operatorname{rank}(T) \geq r$ for every $T \in \mathcal{S}$. If*

$$\operatorname{rank}(L) > |\mathcal{S}| \sum_{j=0}^{\min\{k, n-r\}} \binom{n-r}{j},$$

then there is an $f \in \mathcal{B}_n^{\leq k}$ such that $L(f) \neq 0$ and $fT \equiv 0$ for every $T \in \mathcal{S}$. In particular, if $k \leq (n - r)/3$, it suffices that

$$\operatorname{rank}(L) > 2|\mathcal{S}|\binom{n-r}{k}.$$

Proof. Write $N = \dim \mathcal{B}_n^{\leq k} = \sum_{j=0}^k \binom{n}{j}$. In the basis of multilinear monomials of degree at most k, a polynomial $f \in \mathcal{B}_n^{\leq k}$ is specified by N coefficients. For each $T \in \mathcal{S}$, define $r_T = \operatorname{rank}(T)$. Since the polynomial T is an indicator of a non-empty affine space of codimension r_T, it follows that we can parametrise this space using $n - r_T$ variables, as before:

$$T^{-1}(1) = \{a_T + A_T z : z \in \{0,1\}^{n - r_T}\}.$$

Under this parametrisation, the condition $fT \equiv 0$ is equivalent to $f(a_T + A_T z) \equiv 0$. Affine substitution followed by multilinearisation does not increase degree, so it follows that $f(a_T + A_T z)$ is a multilinear polynomial of degree at most k in $n - r_T$ variables. Therefore, setting its coefficients to zero imposes at most

$$\sum_{j=0}^{\min\{k,n-r_T\}} \binom{n-r_T}{j} \leq \sum_{j=0}^{\min\{k,n-r\}} \binom{n-r}{j}$$

homogeneous linear equations on the coefficients of f.

Let $K \subseteq \mathcal{B}_n^{\leq k}$ be the subspace of polynomials annihilating every $T \in \mathcal{S}$. Combining the equations for all $T \in \mathcal{S}$ gives

$$\dim K \geq N - |\mathcal{S}| \sum_{j=0}^{\min\{k,n-r\}} \binom{n-r}{j} > N - \operatorname{rank}(L) = \dim \ker(L).$$

Thus K is not contained in $\ker(L)$, and any $f \in K \setminus \ker(L)$ has the required properties. For the final assertion, define $m = n - r$. If $k = 0$, then $\sum_{j=0}^{k} \binom{m}{j} = 1 \leq 2\binom{m}{k}$, and otherwise it follows from standard bounds on geometric series. □

By applying the previous lemma, we can find an annihilator for all high-rank terms in a ΣΠΣ-PC proof π. In Section 3, we then converted the resulting proof to a semantic low-degree PC proof, and proved the lower bound directly. In the current abstraction, the pseudo-reduction operators require *syntactic* PC degree lower bounds, which are slightly weaker. Recall from Lemma 2.1 that every semantic PC derivation of degree at most d, with at most two premises per inference, can be simulated by a syntactic PC derivation of degree at most $3d$, from the same axioms and with the same conclusion.

We are now prepared to prove Theorem 4.4.

Proof of Theorem 4.4. Assume for contradiction, that π is a semantic ΣΠΣ-PC refutation of $\mathcal{F}$ such that

$$\operatorname{rank}(\mathcal{R}_{\leq k}) > 2S(\pi)\binom{n-r}{k}.$$

Let $\mathcal{S}$ contain one representative of each distinct non-zero top product term of π having rank at least r. Since $|\mathcal{S}| \leq S(\pi)$ and $3(k+r) \leq B \leq n$ implies $k \leq (n-r)/3$, we may apply Lemma 4.7 with $L = \mathcal{R}_{\leq k}$. This gives an $f \in \mathcal{B}_n^{\leq k}$ such that $\mathcal{R}(f) \neq 0$ and $fT \equiv 0$ for every $T \in \mathcal{S}$.

Write $P_1, \ldots, P_t = 1$ for the lines of π. Each P_i is a sum of top product terms. Multiplication by f annihilates all terms of rank at least r, while every remaining term has degree less than r. Consequently,

$$\deg(fP_i) \leq k + r \qquad \text{for every } i \in [t].$$

We construct a semantic PC derivation of f by replacing each line P_i with the multilinear polynomial representing fP_i. For an axiom line $P_i = G \in \mathcal{F}$, first include G and then infer fG, which is a semantic consequence of G. These two lines have degree at most d_0 and $k + d_0$, respectively.

If P_i follows semantically from P_j and P_ℓ, then fP_i follows semantically from fP_j and fP_ℓ. Indeed, at any Boolean input where $fP_j = fP_\ell = 0$, either $f = 0$, in which case $fP_i = 0$, or $f = 1$, in which case $P_j = P_\ell = 0$ and the original inference gives $P_i = 0$. The final line becomes $fP_t = f$, so this is a semantic derivation of f from $\mathcal{F}$ of degree at most

$$D = k + \max\{d_0, r\}.$$

By Lemma 2.1, there is a syntactic PC derivation of f from $\mathcal{F}$ of degree at most $3D = B$. Since $\mathcal{R}$ is a degree-B pseudo-reduction operator, Lemma 4.2 gives $\mathcal{R}(f) = 0$, contradicting our choice of f. □

4.2 Proof of the $\Sigma\Pi\Sigma\Pi^{[w]}$-PC size bound

We now prove Theorem 4.5, following the structure of the proof of Theorem 4.4. For our applications of this lemma, we will need to use the background axioms, and so before starting the proof we will discuss how to generalise the previous concepts relative to $\mathcal{F}_0$.

Let $\mathcal{F} \subseteq \mathcal{B}_N$ and let $\mathcal{F}_0 \subseteq \mathcal{F}$ be an arbitrary set of background axioms. Recall that

$$\mathcal{B}_{\mathcal{F}_0} = \mathcal{B}_N/\langle \mathcal{F}_0 \rangle$$

is isomorphic to the algebra of Boolean functions on

$$\Omega = \{x \in \{0,1\}^N : G(x) = 0 \text{ for every } G \in \mathcal{F}_0\}.$$

The degree of an element of $\mathcal{B}_{\mathcal{F}_0}$ is the minimum degree of a polynomial representing it. Throughout this section, all identities will be interpreted in $\mathcal{B}_{\mathcal{F}_0}$, and semantic inferences are interpreted on Ω. Passing to this quotient preserves the soundness of a semantic refutation, and can only decrease its number of distinct non-zero top product terms.

We first need to discuss how to generalise the *rank* of a line. Recall that a top product term T in a $\Sigma\Pi\Sigma\Pi^{[w]}$ circuit has the form

$$T = \prod_{j \in J} g_j,$$

where each g_j is a polynomial with degree at most w. Since each g_j is now a polynomial, T is no longer necessarily an indicator of an affine space, and so our parametrisation argument from before no longer applies.

As we mentioned before, we will instead follow the approximation method techniques of Razborov [60] and Smolensky [64]. We will first choose an approximation a_T of degree at most rw for every top-product term T in the proof. Then, instead of choosing a polynomial f which annihilates all high codimension terms, we will choose f so that it annihilates the *error* obtained by replacing T with a_T. Formally speaking, if a polynomial f of degree at most k annihilates the error term $T - a_T$, then

$$f(T - a_T) = 0 \quad \Longrightarrow \quad fT = fa_T,$$

so $\deg(fT) \leq k + rw$. Thus the role of the annihilator is the same as in Lemma 4.7: it makes a specified family of polynomials zero, while remaining outside the kernel of a given linear map.

The following lemma is the analogue of Lemma 4.7. It holds for arbitrary background axioms and, as before, we will apply it with $L = \mathcal{R}_{\leq k}$. Since we are working in $\mathbb{F}_2$, we will write $T + a_T$ instead of $T - a_T$, and note that $(T + a_T)\mathcal{B}_{\mathcal{F}_0}^{\leq k}$ is the space of all multiples $f(T + a_T)$ with $\deg(f) \leq k$ and $f \in \mathcal{B}_{\mathcal{F}_0}$.

Lemma 4.8 (Common annihilator for approximation errors)**.** *Fix integers $1 \leq k \leq N$ and $d \geq 0$, let $\lambda \geq 0$, and let $L \colon \mathcal{B}_{\mathcal{F}_0}^{\leq k} \to W$ be a linear map. Suppose that $\mathcal{S}$ is a finite set of polynomials in $\mathcal{B}_{\mathcal{F}_0}$ and that for every $T \in \mathcal{S}$ there is a polynomial a_T of degree at most d such that*

$$\dim\big((T + a_T)\mathcal{B}_{\mathcal{F}_0}^{\leq k}\big) \leq \lambda.$$

If

$$\operatorname{rank}(L) > |\mathcal{S}|\lambda, \tag{4}$$

then there is an $f \in \mathcal{B}_{\mathcal{F}_0}^{\leq k}$ such that $L(f) \neq 0$ and $f(T + a_T) = 0$ for every $T \in \mathcal{S}$. In particular, $fT = fa_T$ and $\deg(fT) \leq k + d$ for every $T \in \mathcal{S}$.

Proof. Recall that $H_{\mathcal{F}_0}(k) = \dim(\mathcal{B}_{\mathcal{F}_0}^{\leq k})$ and fix a basis of $\mathcal{B}_{\mathcal{F}_0}^{\leq k}$, so that a polynomial $f \in \mathcal{B}_{\mathcal{F}_0}^{\leq k}$ is specified by $H_{\mathcal{F}_0}(k)$ coefficients in this basis. For each $T \in \mathcal{S}$, the product $f(T + a_T)$ depends linearly on these coefficients and belongs to a space of dimension at most λ. Expressing it in a basis of that space shows that the condition $f(T + a_T) = 0$ imposes at most λ homogeneous linear equations on the coefficients of f.

Let $K \subseteq \mathcal{B}_{\mathcal{F}_0}^{\leq k}$ be the subspace of polynomials annihilating every error $T + a_T$. Combining the equations for all $T \in \mathcal{S}$ gives

$$\dim K \geq H_{\mathcal{F}_0}(k) - |\mathcal{S}|\lambda > H_{\mathcal{F}_0}(k) - \operatorname{rank}(L) = \dim \ker(L).$$

Thus K is not contained in $\ker(L)$, and any $f \in K \setminus \ker(L)$ has the required properties. □

It remains to construct approximations for which the spaces of error multiples have small dimension. For this we use the approximation ideas of Razborov [60] and Smolensky [64]. The following statement uses the n-division from Theorem 4.5.

Lemma 4.9 (Approximation of T)**.** *Let $n, h \geq 1$ be integers, put $N = nh$, and partition the variables into $V_1, \ldots, V_n$, each of size h. Let $\mathcal{F}_0$ consist of all products xy of distinct variables in the same part, and work in $\mathcal{B}_{\mathcal{F}_0} = \mathcal{B}_N/\langle \mathcal{F}_0 \rangle$. Let $1 \leq w \leq k \leq n$ and $r \geq 0$ be integers, and let $T = \prod_{j \in J} g_j$, where $\deg(g_j) \leq w$ for each $j \in J$. Then there exist $q_1, \ldots, q_r \in \operatorname{span}\{1 + g_j : j \in J\}$ such that the approximator*

$$a := \prod_{\ell=1}^{r} (1 + q_\ell)$$

has degree at most rw and satisfies $aT = T$. Define the error polynomial $\operatorname{err}_T := T + a$ and the error operator $\mathsf{Err}_T \colon \mathcal{B}_{\mathcal{F}_0}^{\leq k} \to \mathcal{B}_{\mathcal{F}_0}$ by $\mathsf{Err}_T(f) := f\operatorname{err}_T$. Then

$$\operatorname{rank}(\mathsf{Err}_T) = \dim(\operatorname{err}_T \mathcal{B}_{\mathcal{F}_0}^{\leq k}) \leq H_{\mathcal{F}_0}(k)\left(1 - \frac{H_{\mathcal{F}_0}(k-w)}{2H_{\mathcal{F}_0}(k)}\right)^r.$$

If additionally $2w \leq k \leq n/3$, then

$$\operatorname{rank}(\mathsf{Err}_T) \leq \frac{2\binom{n}{k}h^k}{\exp\left(\frac{r}{4}\left(\frac{k}{2N}\right)^w\right)}.$$

The rank bound on the approximation error operator will be used in Section 4.3 to find a common annihilator for the errors introduced when approximating circuit gates. As it is technical, we postpone the proof of Lemma 4.9 until the end of the section. With the previous lemmas in hand, we can prove Theorem 4.5 by essentially the same argument as Theorem 4.4.

First, we need to apply Lemma 4.9 to replace every line of a ΣΠΣΠ-PC proof with a low-degree polynomial, which we can then simulate in semantic polynomial calculus. Luckily, the simulation in Lemma 2.1 extends to $\mathcal{B}_{\mathcal{F}_0}$ with the same degree bound.

Lemma 4.10 (Low-degree simulation of semantic PC by syntactic PC modulo an ideal)**.** *Let $p, q, s \in \mathcal{B}_{\mathcal{F}_0}$ have degree at most d, and suppose that*

$$p(x) = q(x) = 0 \quad \Longrightarrow \quad s(x) = 0 \qquad \text{for every } x \in \Omega.$$

Then s has a PC derivation from p and q in $\mathcal{B}_{\mathcal{F}_0}$ of degree at most $3d$. Consequently, every degree-d semantic PC derivation in $\mathcal{B}_{\mathcal{F}_0}$ with at most two premises per inference can be simulated by a PC derivation in $\mathcal{B}_{\mathcal{F}_0}$ of degree at most $3d$.

Proof. If $p(x) = q(x) = 0$, then $s(x) = 0$ by hypothesis. Otherwise, at least one of $p(x)$ and $q(x)$ is 1. Since we work over $\mathbb{F}_2$, this gives

$$p(x) + q(x) + p(x)q(x) = 1.$$

In either case, multiplying s by $p + q + pq$ does not change its value at x. Thus the following identity holds in $\mathcal{B}_{\mathcal{F}_0}$:

$$s = s(p + q + pq) = s(1 + q)p + sq.$$

The multipliers $s(1 + q)$ and s have degree at most $2d$. Choose representatives of degree at most $2d$ and expand them into monomials. Multiplying p and q by these monomials and adding the results derives s in degree at most $3d$. An inference with one premise is covered by taking $p = q$. Applying this construction to every inference proves the final claim. □

We can now prove Theorem 4.5, following the same argument as in the proof of Theorem 4.4.

Proof of Theorem 4.5. Let π be a semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutation of $\mathcal{F}$, and work in the quotient for the given n-division. Put

$$\lambda = 2\binom{n}{k}h^k \exp\left(-\frac{r}{4}\left(\frac{k}{2N}\right)^w\right).$$

Since $2w \le k$ and $3k \le B \le n$, Lemma 4.9 gives every top product term T of π an approximation a_T of degree at most rw such that

$$\dim\big((T + a_T)\mathcal{B}_{\mathcal{F}_0}^{\le k}\big) \le \lambda.$$

Assume for contradiction that

$$\operatorname{rank}(\mathcal{R}_{\le k}) > S(\pi)\lambda.$$

Let $\mathcal{S}$ contain one representative of each distinct non-zero top product term of π in $\mathcal{B}_{\mathcal{F}_0}$. Since $|\mathcal{S}| = S(\pi)$, we may apply Lemma 4.8 with $L = \mathcal{R}_{\le k}$, $d = rw$, and the above value of λ. This gives an $f \in \mathcal{B}_{\mathcal{F}_0}^{\le k}$ such that $\mathcal{R}(f) \ne 0$ and $f(T + a_T) = 0$ for every $T \in \mathcal{S}$.

Write $P_1, \dots, P_t = 1$ for the lines of π, and note that each P_i is a sum of top product terms. Since $fT = fa_T$ and $\deg(a_T) \le rw$, we have

$$\deg(fP_i) \le k + rw \qquad \text{for every } i \in [t].$$

We construct a semantic PC derivation of f by replacing each line P_i with a polynomial of minimum degree representing fP_i in $\mathcal{B}_{\mathcal{F}_0}$. For an axiom line $P_i = G \in \mathcal{F}$, first include G and then infer fG, which is a semantic consequence of G. These two lines have degree at most d_0 and $k + d_0$, respectively.

If P_i follows semantically from P_j and P_ℓ, then fP_i follows semantically from fP_j and fP_ℓ. Indeed, at any input in Ω where $fP_j = fP_\ell = 0$, either $f = 0$, in which case $fP_i = 0$, or $f = 1$, in which case $P_j = P_\ell = 0$ and the original inference gives $P_i = 0$. The final line becomes $fP_t = f$, so this is a semantic derivation of f from $\mathcal{F}$ of degree at most

$$D = k + \max\{d_0, rw\}.$$

By Lemma 4.10, this implies that there is a syntactic PC derivation of f from $\mathcal{F}$ of degree at most $3D = B$. Since $\mathcal{R}$ is a degree-B pseudo-reduction operator, Lemma 4.2 gives $\mathcal{R}(f) = 0$, contradicting our choice of f. Therefore

$$S(\pi) \ge \frac{\operatorname{rank}(\mathcal{R}_{\le k})}{\lambda} = \frac{\operatorname{rank}(\mathcal{R}_{\le k})}{2\binom{n}{k}h^k} \exp\left(\frac{r}{4}\left(\frac{k}{2N}\right)^w\right). \qquad \square$$

Proof of Corollary 4.6. Apply Theorem 4.5 with $h = 1$. Each part consists of a single variable, so $\mathcal{F}_0 = \varnothing$, $\mathcal{B}_{\mathcal{F}_0} = \mathcal{B}_\varnothing = \mathcal{B}_n$, and $N = n$. $\square$

4.2.1 Proof of the approximation lemma.

We now prove Lemma 4.9. We return to the n-division setting, with $N = nh$ variables and $\mathcal{F}_0$ consisting of the products of distinct variables in the same part.

In the quotient $\mathcal{B}_{\mathcal{F}_0}^{\le k}$, the multilinear monomials containing at most one variable from each part form a basis. Therefore, the dimension function satisfies

$$H_{\mathcal{F}_0}(k) = \dim(\mathcal{B}_{\mathcal{F}_0}^{\le k}) = \sum_{j=0}^{\min\{k,n\}} \binom{n}{j} h^j.$$

Before proving the lemma, we first require the following technical lemma. For any $a \in \mathcal{B}_{\mathcal{F}_0}$ and $j \ge 0$, let

$$r_a(j) = \dim(a\mathcal{B}_{\mathcal{F}_0}^{\le j}).$$

The next lemma shows that the ratio $r_a(j)/H_{\mathcal{F}_0}(j)$ is non-increasing in j.

Lemma 4.11 (Normalised rank monotonicity)**.** *For every $a \in \mathcal{B}_{\mathcal{F}_0}$, the ratio $r_a(j)/H_{\mathcal{F}_0}(j)$ is non-increasing in j. In particular, if $0 \le j \le k \le n$, then*

$$\frac{r_a(j)}{H_{\mathcal{F}_0}(j)} \ge \frac{r_a(k)}{H_{\mathcal{F}_0}(k)}.$$

Proof. Fix a degree-compatible monomial order, and let c_i be the number of degree-i standard monomials[8] for the ideal of polynomials vanishing on $\{x \in \Omega : a(x) = 1\}$. These monomials are closed under taking divisors, contain at most one variable from each part, and satisfy

$$r_a(j) = \sum_{i=0}^{j} c_i.$$

Let $A_i = \binom{n}{i} h^i$ be the number of monomials in $\mathcal{B}_{\mathcal{F}_0}$ of degree i. Every degree-i monomial has i divisors of degree $i-1$, while every degree-$(i-1)$ monomial has $(n-i+1)h$ extensions of degree i, obtained by choosing one of the unused parts and one of its variables. Counting these incidences gives

$$iA_i = (n-i+1)hA_{i-1}.$$

For standard monomials, all these divisors remain standard, whereas some of the extensions may not be standard. The same count therefore gives

$$ic_i \le (n-i+1)hc_{i-1}.$$

Thus the sequence c_i/A_i is non-increasing for $0 \le i \le n$. Each additional term is at most the preceding weighted average, so these averages are also non-increasing for $0 \le j \le n$:

$$\frac{r_a(j)}{H_{\mathcal{F}_0}(j)} = \frac{\sum_{i=0}^{j} c_i}{\sum_{i=0}^{j} A_i}.$$

For $j \ge n$, both $r_a(j)$ and $H_{\mathcal{F}_0}(j)$ are constant, so the claim follows for all $j \ge 0$. □

We can now prove Lemma 4.9.

Proof of Lemma 4.9. We first prove the general rank bound for all integers $1 \le w \le k \le n$ and $r \ge 0$, using $\operatorname{rank}(\mathsf{Err}_T) = \dim((T+a)\mathcal{B}_{\mathcal{F}_0}^{\le k})$. For every product T with factors of degree at most w, we construct

$$q_1, \dots, q_r \in \operatorname{span}\{1 + g_j : j \in J\}$$

such that $a = \prod_{\ell=1}^{r}(1+q_\ell)$ has degree at most rw, satisfies $aT = T$, and obeys

$$\dim((T+a)\mathcal{B}_{\mathcal{F}_0}^{\le k}) \le H_{\mathcal{F}_0}(k)(1-\alpha)^r, \tag{5}$$

where

$$\alpha = \frac{H_{\mathcal{F}_0}(k-w)}{2H_{\mathcal{F}_0}(k)}.$$

If $T = 1$, then we can take $q_\ell = 0$ for every $\ell \in [r]$ and $a = 1$, and the claim is immediate. Otherwise, let $W = \operatorname{span}\{1 + g_j : j \in J\}$. Every $q \in W$ has degree at most w and satisfies $Tq = 0$.

We construct approximations $a_0, \dots, a_r$ to T, with $\deg(a_i) \le iw$, such that $r_{T+a_i}(k)$ decreases by a factor of $1-\alpha$ at each step. This is done in the following claim.

Claim 4.12. For every $b \in \mathcal{B}_{\mathcal{F}_0}$ satisfying $bT = 0$, there is a $q \in W$ such that

$$r_{b(1+q)}(k) \le (1-\alpha) r_b(k).$$

[8] A standard monomial is a monomial that is not divisible by the leading monomial of any polynomial in the ideal.

Proof. Put $V = b\mathcal{B}_{\mathcal{F}_0}^{\leq k-w}$ and $m = \dim(V)$. Choose a basis $v_1, \ldots, v_m$ of V and form the matrix whose rows are the functions $v_1, \ldots, v_m$ evaluated at all inputs in Ω. Since the functions are linearly independent, this matrix has rank m. Choose m independent columns of this matrix, corresponding to inputs $\omega_1, \ldots, \omega_m$.

Choose q uniformly at random from W, and fix any selected input ω_i. Its column is non-zero, and every basis function is a multiple of b, so $b(\omega_i) = 1$ and hence $T(\omega_i) = 0$. Thus some factor g_j vanishes at ω_i, giving $(1 + g_j)(\omega_i) = 1$.

Pair each polynomial $q \in W$ with $q + (1 + g_j)$. Both belong to W, and their values at ω_i differ by 1. These are disjoint pairs, with one member taking the value 0 and the other taking the value 1. Hence exactly half the polynomials in W are 1 at ω_i, so $q(\omega_i) = 1$ with probability $1/2$.

Multiplying the basis functions by q preserves each column when $q = 1$, and sets the other columns to zero. The surviving selected columns remain independent, so linearity of expectation gives

$$\mathbb{E}_{q \in W}\big[\dim(qV)\big] \geq \sum_{i=1}^{m} \Pr_{q \in W}[q(\omega_i) = 1] = \frac{m}{2}.$$

Therefore, there is a $q \in W$ satisfying $\dim(qV) \geq m/2$. Since $qV = bq\mathcal{B}_{\mathcal{F}_0}^{\leq k-w}$, Lemma 4.11 gives

$$r_{bq}(k-w) = \dim(qV) \geq \frac{1}{2} r_b(k-w) \geq \frac{H_{\mathcal{F}_0}(k-w)}{2H_{\mathcal{F}_0}(k)} r_b(k) = \alpha r_b(k). \tag{6}$$

Since $\deg(q) \leq w$, we have $qV \subseteq b\mathcal{B}_{\mathcal{F}_0}^{\leq k}$. Extend a basis of qV to a basis of $b\mathcal{B}_{\mathcal{F}_0}^{\leq k}$. Multiplication by $1 + q$ kills the first $\dim(qV)$ basis vectors, because $q(1 + q) = 0$. The products of the remaining basis vectors with $1 + q$ span $b(1 + q)\mathcal{B}_{\mathcal{F}_0}^{\leq k}$. Thus (6) gives

$$r_{b(1+q)}(k) \leq r_b(k) - \dim(qV) \leq (1 - \alpha) r_b(k),$$

as claimed. ◇

Now, we construct the approximators a_i by induction on $0 \leq i \leq r$ so that $\deg(a_i) \leq iw$. The induction maintains

$$(T + a_i)T = 0$$

together with

$$r_{T+a_i}(k) \leq H_{\mathcal{F}_0}(k)(1 - \alpha)^i.$$

For the base case, set $a_0 = 1$. Since $T^2 = T$ in $\mathcal{B}_{\mathcal{F}_0}$, we have $(T + a_0)T = 0$, while $r_{T+a_0}(k) \leq H_{\mathcal{F}_0}(k)$ because $H_{\mathcal{F}_0}(k) = \dim(\mathcal{B}_{\mathcal{F}_0}^{\leq k})$. Now suppose that $i < r$ and that a_i satisfies the induction hypothesis. Apply the claim with $b = T + a_i$, choose the resulting $q_{i+1} \in W$, and set

$$a_{i+1} = a_i(1 + q_{i+1}).$$

Since $Tq_{i+1} = 0$, we have

$$T + a_{i+1} = (T + a_i)(1 + q_{i+1}).$$

It follows from the induction hypothesis that $(T + a_{i+1})T = 0$, and the claim gives

$$\begin{aligned} r_{T+a_{i+1}}(k) &\leq (1 - \alpha) r_{T+a_i}(k) \\ &\leq H_{\mathcal{F}_0}(k)(1 - \alpha)^{i+1}. \end{aligned}$$

Finally, $\deg(q_{i+1}) \leq w$ gives $\deg(a_{i+1}) \leq (i + 1)w$, completing the induction. Taking $a = a_r = \prod_{\ell=1}^{r}(1 + q_\ell)$ proves (5), together with the required degree bound and identity $aT = T$.

To prove the explicit estimate, assume additionally that $2w \leq k$ and $3k \leq n$. For $1 \leq j \leq k$, the ratio of consecutive summands in $H_{\mathcal{F}_0}(k)$ satisfies

$$\frac{\binom{n}{j-1}h^{j-1}}{\binom{n}{j}h^j} = \frac{j}{h(n - j + 1)} \leq \frac{1}{2}.$$

Summing a geometric series gives $H_{\mathcal{F}_0}(k) \le 2\binom{n}{k}h^k$. On the other hand, $H_{\mathcal{F}_0}(k-w) \ge \binom{n}{k-w}h^{k-w}$, so

$$\alpha \ge \frac{\binom{n}{k-w}}{4\binom{n}{k}h^w} = \frac{1}{4}\prod_{j=0}^{w-1}\frac{k-j}{h(n-k+j+1)} \ge \frac{1}{4}\left(\frac{k}{2N}\right)^w.$$

The last inequality follows from $k-j \ge k/2$ and $h(n-k+j+1) \le nh = N$, since $2w \le k$. Finally, $(1-\alpha)^r \le \exp(-\alpha r)$, so (5) yields

$$\dim((T+a)\mathcal{B}^{\le k}_{\mathcal{F}_0}) \le \frac{2\binom{n}{k}h^k}{\exp\left(\frac{r}{4}\left(\frac{k}{2N}\right)^w\right)}. \qquad \square$$

4.3 Proof of the $\mathsf{AC}^0[2]$-Frege size bound

We approximate circuit arithmetisations by low-degree polynomials by repeatedly applying the approximation from Lemma 4.9 in a bottom-up fashion. The approximation of each AND or OR gate introduces the approximation error err_A. Lemma 4.9 controls the rank of the error operator $\mathsf{Err}_A(f) := f\mathrm{err}_A$. A common annihilator polynomial then annihilates all errors appearing in the proof.

We do not need any background axioms in this section and assume that $\mathcal{F}_0 = \varnothing$, $h = 1$, and $N = n$.

Lemma 4.13 ($\mathsf{AC}^0[2]$-Frege approximation with a common annihilator)**.** *Let $d \ge 2$, and let $\mathcal{F}$ be a family of $\mathsf{AC}^0[2]$ circuits of depth at most d. Let Γ be the set of distinct rooted subcircuits of circuits in $\mathcal{F}$.*

Choose integers $r, R \ge 1$, and set $q = Rr^{d-1}$ and $\alpha_j = \frac{H_n(k-r^{j-1})}{2H_n(k)}$ for each $j \in [d]$. Assume $r^{d-1} \le k \le n$.

There exist polynomial approximations p_A for each $A \in \Gamma$ and a family Err of at most $|\Gamma|$ approximation error polynomials with the following properties.

- *The degree of p_A is at most r^j when A has depth $j < d$, and it is at most q when the depth of A is exactly d.*
- *Let*
$$\mathcal{A} = \{f \in \mathcal{B}^{\le k}_n : f\mathrm{err} = 0 \text{ for every } \mathrm{err} \in \mathsf{Err}\}$$
be the space of all polynomials of degree at most k that annihilate all polynomials in Err. Then
$$\lambda := \dim(\mathcal{A}) \ge (1-\delta)H_n(k), \tag{7}$$
where $\delta = |\Gamma|\exp(-\min\{r\alpha_{d-1}, R\alpha_d\})$.
- *Every $f \in \mathcal{A}$ satisfies $fp_A = f[\![A]\!]$ for every $A \in \Gamma$.*

Proof. Process the circuits in Γ in increasing order of depth. For a literal or constant A, take $p_A = [\![A]\!]$. For an XOR gate with children $A_1, \dots, A_s$, set $p_A = 1 + \sum_{i=1}^s (1 + p_{A_i})$. Consider an OR gate A of depth j with children $A_1, \dots, A_s$. Apply Lemma 4.9 to $\prod_i p_{A_i}$ with degree bound r^{j-1}, using r repetitions if $j < d$ and R repetitions if $j = d$. Let a be the approximation polynomial it constructs, set $p_A = a$, and add the approximation error polynomial

$$\mathrm{err}_A = a + \prod_{i=1}^s p_{A_i}$$

to the set Err. For an AND gate, apply the same lemma to $\prod_i (1 + p_{A_i})$, set $p_A = 1 + a$, and add the approximation error polynomial

$$\mathrm{err}_A = a + \prod_{i=1}^s (1 + p_{A_i}).$$

In both cases, err_A is the difference between the approximation p_A and the gate arithmetisation evaluated on the child approximations.

By induction, every child approximation used at depth j has degree at most r^{j-1}. The lemma gives $\deg(p_A) \le r^j$ for $j < d$ and $\deg(p_A) \le Rr^{d-1} = q$ for $j = d$, and XOR gates preserve these bounds. We only add new errors to Err for AND and OR gates, hence $|\mathsf{Err}| \le |\Gamma|$.

For an AND or OR gate A of depth j, the error operator $\mathsf{Err}_A \colon \mathcal{B}_n^{\le k} \to \mathcal{B}_n$, defined by $\mathsf{Err}_A(f) := f\mathrm{err}_A$, has rank at most $H_n(k)e^{-r\alpha_j}$ for $j < d$ and $H_n(k)e^{-R\alpha_d}$ for $j = d$. The α_j are non-increasing, so each rank is at most $H_n(k)e^{-\theta}$, where $\theta = \min\{r\alpha_{d-1}, R\alpha_d\}$. Since $\mathcal{A}$ is the kernel of the map collecting these errors, rank–nullity gives

$$H_n(k) - \dim(\mathcal{A}) \le \sum_{A \text{ an AND or OR gate}} \operatorname{rank}(\mathsf{Err}_A) \le |\Gamma| H_n(k) e^{-\theta},$$

proving (7).

We prove $fp_A = f[\![A]\!]$ by induction on depth. Fix $f \in \mathcal{A}$ and an input x where $f(x) = 1$. The child approximations equal their arithmetisations at x by induction. At an AND or OR gate A, the condition $f\mathrm{err}_A = 0$ gives $\mathrm{err}_A(x) = 0$, so $p_A(x) = [\![A]\!](x)$. For literals and constants, $p_A = [\![A]\!]$ by definition; for XOR gates, $p_A(x) = [\![A]\!](x)$ follows from the affine formula and the induction hypothesis. Thus $p_A(x) = [\![A]\!](x)$ for every $A \in \Gamma$. At inputs where $f = 0$, both fp_A and $f[\![A]\!]$ vanish. Hence $fp_A = f[\![A]\!]$ in $\mathcal{B}_n$.

□

Here R may be any positive integer: the final product approximation requires only $r^{d-1} \le k$, with no requirement that $q \le k$.

Applying this lemma to an $\mathsf{AC}^0[2]$-Frege refutation gives a PC derivation of a common annihilator f from the arithmetised axioms multiplied by f, with every line of the form fQ for a low-degree polynomial Q.

Lemma 4.14 (Simulation after multiplication)**.** *Let π be a semantic $\mathsf{AC}^0[2]$-Frege refutation of $\mathcal{F}$. Suppose the arithmetisation of each proof line A has a polynomial approximation p_A of degree at most q. If $f \in \mathcal{B}_n^{\le k}$ satisfies $fp_A = f[\![A]\!]$ for every proof line, then there is a PC derivation of f from*

$$\{f[\![A]\!] : A \in \mathcal{F}\},$$

in which every line has the form fQ_i for some $Q_i \in \mathcal{B}_n$ with $\deg Q_i \le 3q$. In particular, the derivation has degree at most $k + 3q$.

Proof. Replace each proof line A in π by fp_A. At inputs where $f = 0$ all these lines vanish. At inputs where $f = 1$, the approximation p_A is equal to $[\![A]\!]$ by assumption. Hence, each inference remains semantically sound. The final constant-false line becomes f.

Suppose that a circuit A is derived from A_1 and A_2 in π. We apply Lemma 2.1 to the semantic PC step

$$\frac{fp_{A_1} \qquad fp_{A_2}}{fp_A}$$

and obtain the polynomial equality

$$fp_A = p_A(1 + p_{A_2})(fp_{A_1}) + p_A(fp_{A_2}). \tag{8}$$

Here $f^2 = f$ removes the extra factors of f in the simulation identity, leaving multipliers of degree at most $2q$ and q, respectively. Since p_A, p_{A_1}, and p_{A_2} have degree at most q, this step can be simulated in PC in degree at most $k + 3q$ by a derivation in which every line has a representation fQ with $\deg(Q) \le 3q$.

Thus, we constructed a syntactic PC derivation of f that is contained entirely in $f\mathcal{B}_n^{\le 3q}$. □

Observe that PC derivations produced by Lemma 4.14 have a specific form: every proof line in the constructed f derivation has the form fQ for some Q with $\deg(Q) \le 3q$. We generalise pseudo-reduction operators from Definition 4.1 relative to derivations of f having this form.

Definition 4.15 (f-relative pseudo-reduction operator)**.** Let $\mathcal{S} \subseteq \mathcal{B}_n$, let $0 \neq f \in \mathcal{B}_n$, and let $k \geq 0$ be an integer. A *degree-k f-relative pseudo-reduction for $\mathcal{S}$* is a linear map

$$\mathcal{R}_f \colon f\mathcal{B}_n^{\leq k} \to W$$

such that

- $\mathcal{R}_f(fa) = 0$ for every $a \in \mathcal{S}$ with $fa \in f\mathcal{B}_n^{\leq k}$.
- For every variable x and every $u \in f\mathcal{B}_n^{\leq k}$ with $xu \in f\mathcal{B}_n^{\leq k}$, if $\mathcal{R}_f(u) = 0$ then $\mathcal{R}_f(xu) = 0$.

Every such operator annihilates all lines of a PC derivation from $f\mathcal{S}$ that stays in $f\mathcal{B}_n^{\leq k}$. Moreover, $\mathcal{R}_f$ is non-zero if and only if $\mathcal{R}_f(f) \neq 0$. Indeed, if $\mathcal{R}_f(f) = 0$, successive variable multiplications show that it vanishes on every f times a monomial of degree at most k, and these products span its domain.

There is a naïve relative pseudo-reduction, analogous to the one from Proposition 4.3. For an integer $\ell \geq 0$, let $D_\ell^{(f)}(\mathcal{S})$ be the space of polynomials admitting a PC derivation from $f\mathcal{S}$ entirely in $f\mathcal{B}_n^{\leq \ell}$. Addition of derivations shows that this is a linear subspace, and appending a permitted multiplication step shows that it has the required closure property. Thus the quotient map

$$\mathcal{R}_{f,\ell} \colon f\mathcal{B}_n^{\leq \ell} \to f\mathcal{B}_n^{\leq \ell} / D_\ell^{(f)}(\mathcal{S}) \tag{9}$$

is an f-relative pseudo-reduction. It is non-zero precisely when f has no such derivation.

Theorem 4.16 (Size bound via relative pseudo-reduction operators)**.** *Let $\mathcal{F}$ be a family of circuits in n variables, and put $\mathcal{S} = \{[\![A]\!] : A \in \mathcal{F}\}$. Fix $d \geq 2$, integers $r, R \geq 1$ and $r^{d-1} \leq k \leq n$, and put*

$$q = Rr^{d-1}, \qquad \alpha_j = \frac{H_n(k - r^{j-1})}{2H_n(k)}, \qquad \theta = \min\{r\alpha_{d-1}, R\alpha_d\}.$$

Suppose $U \subseteq \mathcal{B}_n^{\leq k}$ is a linear subspace such that, for every $0 \neq f \in U$, there is a non-zero f-relative pseudo-reduction of degree $3q$ for $\mathcal{S}$. Then every depth-d semantic $\mathsf{AC}^0[2]$-Frege refutation π of $\mathcal{F}$ satisfies

$$\operatorname{size}(\pi) \geq \frac{\dim(U)}{H_n(k)} \exp(\theta). \tag{10}$$

Proof. Let Γ be the set of distinct rooted subcircuits in π, so $|\Gamma| = \operatorname{size}(\pi)$. By Lemma 4.13, the common error-annihilating space $\mathcal{A} \subseteq \mathcal{B}_n^{\leq k}$ has codimension at most $H_n(k)|\Gamma|e^{-\theta}$. Hence

$$\dim(U \cap \mathcal{A}) \geq \dim(U) - H_n(k)|\Gamma|e^{-\theta}.$$

If $H_n(k)|\Gamma|e^{-\theta} < \dim(U)$, there would be a non-zero $f \in U \cap \mathcal{A}$. By Lemma 4.14, f would have a PC derivation from $f\mathcal{S}$ entirely in $f\mathcal{B}_n^{\leq 3q}$. Its f-relative pseudo-reduction would therefore satisfy $\mathcal{R}_f(f) = 0$, contradicting that it is non-zero. Hence $\dim(U) \leq H_n(k)|\Gamma|e^{-\theta}$, proving the claim. □

We now make the dependence of the approximation loss on the degree cutoff explicit.

Lemma 4.17. *Let $0 < \epsilon \leq 1/256$, $n \geq 16/\epsilon^2$, and*

$$k = \lfloor (1/2 - \epsilon)n \rfloor.$$

For every integer $0 \leq a \leq \epsilon n/2$,

$$\frac{H_n(k-a)}{H_n(k)} \geq e^{-128\epsilon a}. \tag{11}$$

Proof. Put $t = \lfloor 1/(32\epsilon) \rfloor$. For integers $k - \epsilon n/2 \leq j \leq k$ and $0 \leq s \leq t$, the assumptions imply that the indices below are nonnegative and

$$\frac{\binom{n}{j-s-1}}{\binom{n}{j-s}} = \frac{j-s}{n-j+s+1} \geq \frac{1/2 - 2\epsilon}{1/2 + 2\epsilon} \geq 1 - 8\epsilon.$$

Consequently, $\binom{n}{j-s} \geq (1-8\epsilon)^s \binom{n}{j} \geq \frac{3}{4}\binom{n}{j}$. Summing over $0 \leq s \leq t$ gives

$$H_n(j) \geq \frac{1}{64\epsilon}\binom{n}{j}, \qquad \frac{H_n(j-1)}{H_n(j)} = 1 - \frac{\binom{n}{j}}{H_n(j)} \geq 1 - 64\epsilon \geq e^{-128\epsilon}.$$

Hence

$$\frac{H_n(k-a)}{H_n(k)} = \prod_{j=k-a+1}^{k} \frac{H_n(j-1)}{H_n(j)} \geq e^{-128\epsilon a}.$$

□

Combining Theorem 4.16 and Lemma 4.17 gives the following size bound.

Corollary 4.18 (Size bound for semantic $\mathsf{AC}^0[2]$-Frege)**.** *Let $\mathcal{F}$ be a family of $\mathsf{AC}^0[2]$ circuits in n variables, and let $\mathcal{S} = \{[\![A]\!] : A \in \mathcal{F}\}$ be its arithmetisation. Fix integers $d \geq 2$ and $r, R \geq 1$, and a parameter $0 < \epsilon \leq 1/256$, and put*

$$\begin{aligned} k &:= \lfloor (1/2 - \epsilon)n \rfloor, \\ q &:= Rr^{d-1}, \\ \theta &:= \frac{1}{2}\min\left\{ re^{-128\epsilon r^{d-2}}, Re^{-128\epsilon r^{d-1}} \right\}. \end{aligned}$$

Suppose that $n \geq 16/\epsilon^2$, $r^{d-1} \leq \epsilon n/2$, and that $U \subseteq \mathcal{B}_n^{\leq k}$ is a linear subspace such that, for every $0 \neq f \in U$, there is a non-zero degree-$3q$ f-relative pseudo-reduction operator for $\mathcal{S}$. Then every semantic depth-d $\mathsf{AC}^0[2]$-Frege refutation π of $\mathcal{F}$ satisfies

$$\operatorname{size}(\pi) \geq \frac{\dim(U)}{H_n(k)} \exp(\theta). \tag{12}$$

Proof. The assumptions imply $r^{d-1} \leq k \leq n$. Applying Lemma 4.17 with $a = r^{j-1}$ for each $j \in [d]$ gives

$$\frac{H_n(k - r^{j-1})}{2H_n(k)} \geq \frac{1}{2} e^{-128\epsilon r^{j-1}}.$$

Substituting these estimates into the bound of Theorem 4.16 proves the claim. □

5 Rank Bounds for Existing Pseudo-Reduction Operators

Here we introduce the expansion pseudo-reduction operator used in the random-formula and graph pigeonhole applications. The construction is a special case of the method of Mikša and Nordström [54], which generalises the Alekhnovich–Razborov method [2]. We then give two rank estimates: one shared by logarithmic-width CNFs and graph PHP, and one for bounded-occurrence 2- and 3-CNFs.

5.1 The expansion pseudo-reduction operator

Let $G = (L \sqcup R, E)$ be a bipartite graph. For $J \subseteq L$, let $N_G(J)$ be its neighbourhood and let $\partial_G(J)$ be its boundary, i.e., the set of unique neighbours of J:

$$\partial_G(J) = \{v \in R : |N_G(v) \cap J| = 1\}.$$

We call G an (s, c)-expander if $|N_G(J)| \geq c|J|$ for every $J \subseteq L$ with $|J| \leq s$, and an (s, c)-boundary expander if $|\partial_G(J)| \geq c|J|$ for all such J.

Suppose that we have an unsatisfiable CNF formula with polynomial encoding $\mathcal{F}$. To apply the graph terminology above, write

$$\mathcal{F} = \mathcal{F}_0 \cup \bigcup_{u \in L} \mathcal{F}_u,$$

where $\mathcal{F}_0$ is a set of background axioms and the sets $\mathcal{F}_u$ are the left axiom clusters. We assume that the background axioms have a common Boolean zero. Partition the polynomial variables into right variable clusters $\{V_v : v \in R\}$, and join u and v whenever a polynomial in $\mathcal{F}_u$ contains a variable from V_v.

The clustered representation is *respectful* if, for every $J \subseteq L$ and every $v \in R$ having a unique neighbour u in J, there is an assignment to V_v that sets every polynomial in $\mathcal{F}_u$ to zero and either sets to zero or leaves unchanged each background axiom in $\mathcal{F}_0$. This is the algebraic form of the respectful-boundary condition in [54]. We describe a general respectful representation for any CNF formula and a more specific one for the graph pigeonhole principle used in Section 7.2:

- For a CNF formula, $\mathcal{F}_0 = \varnothing$, the left clusters are individual clause polynomials, and $V_i = \{x_i\}$. A boundary variable can be assigned to satisfy its unique neighbouring clause.
- For the graph pigeonhole principle, defined in Section 7.2, the background axioms prohibit two pigeons from occupying the same hole, the left clusters assert that individual pigeons occupy adjacent holes, and
$$V_v = \{x_{u,v} : u \in N_G(v)\}$$
is the variable cluster associated with a hole v. If v is a boundary hole for a set of pigeons and u is its unique neighbour, set $x_{u,v} = 1$ and all other variables in V_v to zero. This satisfies the axiom for u and all background axioms involving variables in V_v.

The operator reduces each monomial using only a small collection of axiom clusters. To specify this collection, for $A \subseteq R$ put

$$\mathrm{Sup}_s(A) = \bigcup\{J \subseteq L : |J| \le s \text{ and } \partial_G(J) \subseteq A\}.$$

Thus the *s-support* of A collects the small sets of left vertices whose boundary is entirely covered by A. For a monomial t, let

$$N(t) = \{v \in R : t \text{ contains a variable from } V_v\}, \qquad \mathrm{Sup}_s(t) = \mathrm{Sup}_s(N(t)).$$

Proposition 5.1 (Small supports)**.** *Let G be a finite (s, c)-boundary expander with $c > 0$. If $|A| \le cs/2$, then*

$$\partial_G(\mathrm{Sup}_s(A)) \subseteq A, \qquad |\mathrm{Sup}_s(A)| \le |A|/c \le s/2.$$

Proof. Every set J in the defining union satisfies $c|J| \le |\partial_G(J)| \le |A|$, so $|J| \le s/2$. If J_1 and J_2 occur in this union, then $|J_1 \cup J_2| \le s$ and

$$\partial_G(J_1 \cup J_2) \subseteq \partial_G(J_1) \cup \partial_G(J_2) \subseteq A.$$

Hence their union also occurs. Since G is finite, the full union $\mathrm{Sup}_s(A)$ satisfies the same conditions and the same size bound. □

Fix a degree-compatible monomial order. For an ideal I in the Boolean polynomial ring $\mathcal{B}$, call a monomial *irreducible* modulo I if it is not the leading monomial of a polynomial in I. Write $R_I(p)$ for the unique linear combination of irreducible monomials such that $p - R_I(p) \in I$. This normal form does not increase degree. For $J \subseteq L$, put

$$I(J) = \left\langle \mathcal{F}_0 \cup \bigcup_{u \in J} \mathcal{F}_u \right\rangle.$$

As in the abstract framework, we work modulo the background axioms:

$$\mathcal{B}_{\mathcal{F}_0} = \mathcal{B}/\langle \mathcal{F}_0 \rangle.$$

Its standard monomial basis consists of the monomials irreducible modulo $\langle \mathcal{F}_0 \rangle$. Since this ideal is contained in every $I(J)$, normal forms modulo $I(J)$ also lie in this basis. When $\mathcal{F}_0 = \varnothing$, the quotient is simply $\mathcal{B}$ and every multilinear monomial is standard.

Definition 5.2 (Expansion pseudo-reduction operator)**.** Suppose that the clustered representation is respectful and G is an (s, c)-boundary expander. Fix an integer K such that every axiom has degree at most K and $K \le cs/2 - 1$. For each standard monomial t of degree at most K, define

$$\mathcal{R}(t) = R_{I(\mathrm{Sup}_s(t))}(t),$$

and extend $\mathcal{R}$ linearly to $\mathcal{B}_{\mathcal{F}_0}^{\le K}$.

By Proposition 5.1, each local reduction uses axioms from at most $s/2$ left clusters, in addition to the background axioms. The following is the key property of this construction.

Lemma 5.3 (Mikša–Nordström reduction [54])**.** *The operator in Definition 5.2 does not increase degree, satisfies $\mathcal{R}(f) = 0$ for every $f \in \mathcal{F}$, and obeys*

$$\mathcal{R}(p) = 0 \quad \Longrightarrow \quad \mathcal{R}(xp) = 0$$

for every polynomial variable x whenever $\deg(xp) \le K$. Thus it is a degree-K pseudo-reduction operator for $\mathcal{F}$.

We will estimate its rank by counting fixed monomials. If a standard monomial t has $\mathrm{Sup}_s(t) = \varnothing$, then directly from the definition,

$$\mathcal{R}(t) = R_{I(\varnothing)}(t) = R_{\langle \mathcal{F}_0 \rangle}(t) = t. \tag{13}$$

Boundary expansion gives $\mathrm{Sup}_s(1) = \varnothing$, and the background axioms have a common Boolean zero, so 1 is standard and $\mathcal{R}(1) = 1$. Distinct fixed standard monomials are linearly independent, so their number is a lower bound on the rank of $\mathcal{R}$.

5.2 Rank estimates

Call $A \subseteq R$ *s-boundary-avoiding* if it contains no boundary $\partial_G(J)$ with $0 < |J| \le s$. Equivalently, $\mathrm{Sup}_s(A) = \varnothing$. Thus every standard monomial of degree at most K using only clusters indexed by A is fixed by $\mathcal{R}$.

The first estimate counts boundary-avoiding sets using connected sets of left vertices. It will be applied to both logarithmic-width CNFs and PHP.

Theorem 5.4 (Rank from boundary expansion)**.** *Let $G = (L \sqcup R, E)$ be an (s, c)-boundary expander, and let $\mathcal{R}$ be the associated operator from Definition 5.2, defined up to degree K. Write $m = |L|$ and $n = |R|$. Suppose that every right cluster has the same positive size D and every monomial containing at most one variable from each cluster is irreducible modulo $\langle \mathcal{F}_0 \rangle$. Let $\kappa \ge 1$ bound the maximum degree of the graph on L in which two vertices are adjacent if they share a right neighbour. For $1 \le k \le \min\{K, n\}$ put $p = k/n$. If $4\kappa p^c < 1$, then*

$$\mathrm{rank}(\mathcal{R}_{\le k}) \ge \left(1 - \frac{mp^c}{1 - 4\kappa p^c}\right) \binom{n}{k} D^k. \tag{14}$$

If the left and right degrees of G are at most Δ_L and Δ_R, respectively, one may take $\kappa = \max\{1, \Delta_L(\Delta_R - 1)\}$.

Proof. Let H be the graph on L in the statement, and choose a k-element set $A \subseteq R$ uniformly at random. If $\partial_G(J) \subseteq A$ for some $0 < |J| \le s$, the same holds for a connected component of $H[J]$: distinct components have disjoint neighbourhoods in G, so the boundary of each component is contained in $\partial_G(J)$.

There are at most $m(4\kappa)^{j-1}$ connected j-vertex sets in H. Indeed, encode such a set by a canonical rooted ordered spanning tree. There are at most 4^{j-1} tree shapes, m choices for the root, and at most κ choices for the image of each remaining vertex given its parent. For any fixed J with $|J| = j \le s$,

$$\Pr[\partial_G(J) \subseteq A] \le p^{|\partial_G(J)|} \le p^{cj}.$$

The first inequality follows by sampling A without replacement; if $|\partial_G(J)| > k$, the probability is zero. A union bound gives

$$\begin{aligned}\Pr[A \text{ is not } s\text{-boundary-avoiding}] &\le \sum_{j=1}^{\lfloor s \rfloor} m(4\kappa)^{j-1} p^{cj} \\ &\le \frac{mp^c}{1-4\kappa p^c}.\end{aligned}$$

Each boundary-avoiding k-set yields D^k distinct standard monomials, by choosing one variable from each of its clusters. All are fixed by (13), and monomials from different choices are linearly independent. Counting them proves the bound. □

For bounded-occurrence 2- and 3-CNFs, we instead count monomials whose variables are far apart in the incidence graph.

Lemma 5.5 (Rank for bounded-occurrence 2- and 3-CNFs)**.** *Let F be a CNF on n variables whose clauses have width two or three and whose variables occur in at most d clauses. Suppose its incidence graph G is an (s, c)-boundary expander with $c > 1/2$, and let $\mathcal{R}$ be its expansion operator, with $\mathcal{F}_0 = \varnothing$, defined up to degree K. Put*

$$b = 1 + 2d + (2d)^2 + (2d)^3.$$

For $1 \le k \le K$ with $b(k-1) < n$, we have

$$\operatorname{rank}(\mathcal{R}_{\le k}) \ge \frac{(n - b(k-1))^k}{k!}.$$

Proof. Form the primal graph on the variables, joining two variables if they occur in a common clause. Its maximum degree is at most $2d$, so every ball of radius three contains at most b vertices.

Every set A of variables at pairwise distance at least four in this graph is s-boundary-avoiding. Indeed, suppose that $B = \partial_G(J) \subseteq A$ for some $0 < |J| \le s$. For each $v \in B$, let $u_v \in J$ be its unique neighbouring clause. These clauses are distinct, since otherwise two vertices of B would be adjacent. Moreover, u_v shares a variable with another clause $w_v \in J$; otherwise all its variables would lie in B, contradicting separation, since u_v has at least two variables. The pairs $\{u_v, w_v\}$ are disjoint: any overlap would put the corresponding boundary vertices at distance at most three. Hence $2|B| \le |J|$, contradicting $|B| \ge c|J| > |J|/2$.

Choose k variables sequentially, maintaining pairwise distance at least four. Each choice excludes at most b variables, so there are at least $(n - b(k-1))^k$ ordered choices. Dividing by $k!$ counts distinct degree-k monomials. They are fixed by (13) and are linearly independent, proving the rank bound. □

6 Lower Bounds for Semantic ΣΠΣ-PC

This section proves exponential lower bounds for semantic ΣΠΣ-PC refutations over $\mathbb{F}_2$, allowing any semantically sound inference from at most two premises, as defined in the preliminaries.

In Section 6.1, we prove an exponential lower bound for random 3-CNFs with a linear number of clauses. We first bound the variable occurrences while preserving boundary expansion and then apply the rank estimate for the Mikša–Nordström operator from Section 5.

In Section 6.2, we prove an exponential lower bound for the functional pigeonhole principle on bounded-degree boundary expanders with a planted matching. This yields an exponential lower bound for the complete functional pigeonhole principle and improves the lower bound from Section 3 for the bit pigeonhole principle with $n + 1$ pigeons and n holes.

Finally, Section 6.3 gives polynomial-size semantic $\Sigma\Pi\Sigma\Pi^{[3]}$-PC refutations of random 3-CNFs and semantic $\Sigma\Pi\Sigma\Pi^{[2]}$-PC refutations of the functional pigeonhole principle. Together with the preceding lower bounds, these give exponential separations from semantic ΣΠΣ-PC.

6.1 Random 3-CNFs

Fix a constant $\alpha > 0$ and let $m = \lfloor \alpha n \rfloor$. Define the random 3-CNF formula $F_{n,\alpha}$ on variables $x_1, \dots, x_n$ by

$$F_{n,\alpha} = \bigwedge_{j=1}^{m} C_j,$$

where the clauses are sampled independently as follows. For every $j \in [m]$, choose a clause C_j from the set of all $8\binom{n}{3}$ clauses of width exactly three uniformly at random. The clauses are sampled with repetition, so the same clause may appear more than once. When n and α are understood, we write $F = F_{n,\alpha}$ and identify F with its polynomial encoding $\mathcal{F}_F = \{[\![C_j]\!] : j \in [m]\}$ from the preliminaries.

Friedgut [31] proved that random 3-CNF formulas have a sharp satisfiability threshold c_n in their clause density; we state below the simpler upper bound due to Franco and Paull [30] that is sufficient for our purposes.

Lemma 6.1 (Unsatisfiability of random 3-CNFs [30])**.** *If*

$$\alpha > \frac{\ln 2}{\ln(8/7)},$$

then $F_{n,\alpha}$ is unsatisfiable with probability $1 - o(1)$.

Proof. The expected number of satisfying assignments is $2^n(7/8)^m = o(1)$. The claim follows from Markov's inequality. □

The *clause–variable incidence graph* of F is the bipartite graph

$$G(F) = ([m] \sqcup [n], E_F),$$

where $j \in [m]$ is adjacent to $i \in [n]$ if the variable x_i occurs in C_j, with either sign. Thus every clause vertex has degree three. We write $N_F = N_{G(F)}$ and $\partial_F = \partial_{G(F)}$.

Lemma 6.2 (Expansion of random 3-CNFs [22, 1])**.** *For every fixed $\alpha > 0$, there is a constant $0 < \gamma = \gamma(\alpha) \le 1$ such that, with probability $1 - o(1)$, $G(F_{n,\alpha})$ is a $(\gamma n, 15/8)$-expander. Consequently, it is a $(\gamma n, 3/4)$-boundary expander.*

The Mikša–Nordström operator already applies to this incidence graph, but its rank need not be large when a variable occurs in many clauses. We therefore first apply Lemma 6.3. The deleted clauses can be satisfied by peeling along their boundary variables.

The following cleaning lemma gives a subgraph with bounded right degree while preserving expansion. Related degree-reduction arguments appear in Razborov's work on polynomial calculus lower bounds [59]; the lemma below is a variant of the cleaning arguments used for random formulas [1]. The proof is in Appendix B.

Lemma 6.3 (Right-degree reduction)**.** *Let $G = (L \sqcup R, E)$ be a bipartite graph with $m \le \alpha n$ left vertices and n right vertices, and every left vertex of degree three. Suppose that $0 < \gamma \le 1$ and that G is a $(\gamma n, 15/8)$-expander. Then, for all sufficiently large n, there are sets $I \subseteq L$ and $A \subseteq R$ such that*

$$|I| \le \frac{\gamma n}{4},$$
$$|I| < |A| < \frac{n}{2},$$
$$N_G(I) \subseteq A.$$

If

$$G' = ((L \setminus I) \sqcup (R \setminus A), E'),$$
$$E' = E \cap ((L \setminus I) \times (R \setminus A)),$$

then:

1. *every left vertex of G' has degree two or three;*

2. *every right vertex of G' has degree at most some positive integer $d = d(\alpha, \gamma)$;*

3. *G' is a $(\gamma n/2, 7/4)$-expander;*

4. *$|\partial_{G'}(J)| > |J|/2$ for every non-empty $J \subseteq L \setminus I$ with $|J| \leq \gamma n/2$.*

Corollary 6.4 (Right-degree reduction for random 3-CNFs)**.** *For every fixed $\alpha > 0$, there are a constant $\gamma = \gamma(\alpha) > 0$ and a positive integer $d = d(\alpha)$ such that, with probability $1 - o(1)$, there is a partial assignment ρ assigning at most $n/2$ variables such that in $F' = F_{n,\alpha} \restriction \rho$ every clause has width two or three, every variable occurs in at most d clauses, and*

$$|\partial_{F'}(J)| > \frac{1}{2}|J|$$

for every non-empty set J of at most $\gamma n/2$ clauses of F'.

Proof. Suppose that $G(F_{n,\alpha})$ satisfies the conclusions of Lemma 6.2. Apply Lemma 6.3, and let I and A be the resulting sets of clauses and variables.

Every non-empty $J \subseteq I$ has a boundary variable, since

$$|\partial_F(J)| \geq 2|N_F(J)| - 3|J| \geq \frac{3}{4}|J| > 0.$$

Hence the clauses in I can be peeled and satisfied on A by assigning the chosen boundary variables in reverse order. Extend this to an assignment ρ of all variables in A. The clauses of $F' = F_{n,\alpha} \restriction \rho$ form a subgraph of the graph G' in Lemma 6.3, so the claimed degree and boundary properties follow. □

We now apply the common framework from Section 5. There are no background axioms, every clause polynomial is a left axiom cluster, and every variable is a singleton right cluster. Thus the clustered representation is respectful, and every monomial with empty support is irreducible modulo the background ideal. The bounded variable occurrence supplied by Corollary 6.4 therefore turns Lemma 5.5 into the required rank estimate.

Theorem 6.5 (Random 3-CNFs)**.** *For every fixed $\alpha > 0$, with probability $1 - o(1)$ every semantic $\Sigma\Pi\Sigma$-PC refutation π of $F_{n,\alpha}$ satisfies*

$$S(\pi) \geq 2^{\Omega_\alpha(n)}.$$

Consequently, its $\Sigma\Pi\Sigma$-size is at least $2^{\Omega_\alpha(n)}$. If $\alpha > \ln 2/\ln(8/7)$, then $F_{n,\alpha}$ is also unsatisfiable with probability $1 - o(1)$.

Proof. By Corollary 6.4, with probability $1-o(1)$ there is a restriction ρ assigning at most $n/2$ variables such that $F' = F_{n,\alpha} \restriction \rho$ has the stated boundary and degree properties. Fix such a restriction and relabel the remaining variables by $[N]$, where $n/2 \leq N \leq n$, including variables that occur in no clause.

Put

$$s = \left\lfloor \frac{\gamma n}{2} \right\rfloor, \qquad b = 1 + 2d + (2d)^2 + (2d)^3.$$

The incidence graph $G(F')$ has left degrees two or three and right degree at most d. Its strict boundary condition implies that it is an (s, c)-boundary expander for some $c > 1/2$.

Choose an integer $k = \Theta_\alpha(n)$ sufficiently small that, with

$$r = (b+1)k, \qquad B = 3(k+r) = 3(b+2)k,$$

we have

$$B \leq \min\left\{N, \frac{s}{4} - 1\right\}.$$

For sufficiently large n, we also have $r \geq 3$.

Let

$$\mathcal{R} \colon \mathcal{B}_N^{\leq B} \to \mathcal{B}_N^{\leq B}$$

be the expansion operator for F' from Definition 5.2. Since $B \le s/4 - 1 < cs/2 - 1$, Lemma 5.3 shows that $\mathcal{R}$ is a degree-B pseudo-reduction operator.

Since F' has clause widths two or three and variable occurrences at most d, Lemma 5.5 applies. As $b(k-1) < N$, it gives

$$\operatorname{rank}(\mathcal{R}_{\le k}) \ge \frac{(N - b(k-1))^k}{k!}.$$

Restricting a semantic $\Sigma\Pi\Sigma$-PC refutation π of $F_{n,\alpha}$ by ρ yields a semantic $\Sigma\Pi\Sigma$-PC refutation π' of F' with $S(\pi') \le S(\pi)$. Since the axioms of F' have degree at most three and $r \ge 3$, Theorem 4.4 gives

$$\begin{aligned} S(\pi) \ge S(\pi') &\ge \frac{\operatorname{rank}(\mathcal{R}_{\le k})}{2\binom{N-r}{k}} \\ &\ge \frac{1}{2}\left(\frac{N - b(k-1)}{N-r}\right)^k \\ &= \frac{1}{2}\left(1 + \frac{k+b}{N-r}\right)^k \\ &\ge \frac{1}{2}\left(1 + \frac{k}{N}\right)^k = 2^{\Omega_\alpha(n)}. \end{aligned}$$

Here we used $\binom{N-r}{k} \le (N-r)^k/k!$, $k = \Theta_\alpha(n)$, and $n/2 \le N \le n$. The same lower bound holds for the $\Sigma\Pi\Sigma$-size, which is at least $S(\pi)$.

The unsatisfiability assertion follows from Lemma 6.1. □

6.2 Functional pigeonhole principle

Definition 6.6 (Functional pigeonhole principle)**.** Let $m > n$, and let $x_{u,v}$ denote that pigeon $u \in [m]$ is mapped to hole $v \in [n]$. The polynomial encoding of the functional pigeonhole principle FPHP_n^m over $\mathbb{F}_2$ consists of the pigeon axioms

$$P_u := \prod_{v \in [n]} (1 + x_{u,v}) \qquad (u \in [m]), \tag{15}$$

the functionality axioms

$$F_{u;v,v'} := x_{u,v}x_{u,v'} \qquad (u \in [m],\ v \ne v'), \tag{16}$$

and the hole axioms

$$H_{u,u';v} := x_{u,v}x_{u',v} \qquad (u \ne u',\ v \in [n]). \tag{17}$$

For a bipartite graph $G = (L \sqcup R, E)$ with $|L| > |R|$, the graph functional pigeonhole principle $\mathsf{FPHP}(G)$ has variables $x_{u,v}$ for $(u,v) \in E$. Its axioms are those in (15)–(17), with pigeon products taken over $N_G(u)$ and quadratic axioms included only when all the corresponding edges belong to E. Thus $\mathsf{FPHP}_n^m = \mathsf{FPHP}(K_{m,n})$.

In this subsection we prove that every semantic $\Sigma\Pi\Sigma$-PC refutation of FPHP_n^{n+1} has size $2^{\Omega(n)}$. It suffices to prove this lower bound for the restriction $\mathsf{FPHP}(G + M)$, where G is a boundary expander with $n+1$ pigeons and n holes and M is a planted matching covering all holes. We establish this bound using Theorem 4.4.

Let $G = (L \sqcup R, E)$ be an (s, δ_0)-boundary expander with $n+1$ pigeons and n holes, where $s \ge 1$ and $\delta_0 > 3$. Assume that its left and right degrees are at most the positive integers $\Delta_{L,0}$ and $\Delta_{R,0}$, respectively. Choose $u_* \in L$, put $L_0 = L \setminus \{u_*\}$, and fix a bijection $p\colon L_0 \to R$. Define

$$M = \{(u, p(u)) : u \in L_0\}$$

and let $H = G + M = (L \sqcup R, E \cup M)$. Put

$$\begin{aligned} \Delta_L &= \Delta_{L,0} + 1, \\ \Delta_R &= \Delta_{R,0} + 1, \\ \delta &= \delta_0 - 1 > 2. \end{aligned}$$

The graph H has left degree at most Δ_L and right degree at most Δ_R. It is also an (s, δ)-boundary expander: adding M introduces at most $|J|$ edges incident with any set $J \subseteq L$, so at most $|J|$ vertices can be lost from its boundary. The existence of suitable graphs G is proved in Lemma 6.7.

The following existence result follows from the work of Ben-Sasson, Harsha, and Raskhodnikova.

Lemma 6.7 ([13])**.** *There exist positive integers $\Delta_{L,0}, \Delta_{R,0}$ and a constant $\gamma > 0$ such that, for all sufficiently large n, there is a $(\gamma n, 4)$-boundary expander with $n+1$ pigeons, n holes, left degree at most $\Delta_{L,0}$, and right degree at most $\Delta_{R,0}$.*

Write $\mathcal{F}$ for the polynomial encoding of $\mathsf{FPHP}(H)$ over $\mathbb{F}_2$. Its axioms are those in (15)–(17), with pigeon products taken over $N_H(u)$ and quadratic axioms included only for edges of H.

We first apply an affine substitution σ to the matching variables. Variables on edges outside M are left unchanged. We consider σ such that

$$\sigma(x_{u,p(u)}) = 1 + \sum_{v \in N_H(u) \setminus \{p(u)\}} x_{u,v} \qquad (u \in L_0). \tag{18}$$

Put $\mathcal{F}' = \sigma(\mathcal{F})$ and

$$n_{\text{free}} = |E(H)| - n,$$
$$d_* = |N_H(u_*)|.$$

We call the n_{free} variables outside M the *free variables*. Of these, $d_* \le \Delta_L$ belong to the unmatched pigeon u_*.

Lemma 6.8 (Affine elimination)**.** *The axioms of $\mathcal{F}'$ have degree at most Δ_L. For $B \ge \Delta_L + 1$,*

$$D_B(\mathcal{F}') \subseteq D_B(\mathcal{F}), \tag{19}$$

where polynomials in the free variables are identified with the same polynomials in the original variables. Moreover, σ maps every semantic $\Sigma\Pi\Sigma$-PC refutation π of $\mathcal{F}$ to a semantic $\Sigma\Pi\Sigma$-PC refutation π' of $\mathcal{F}'$ with $S(\pi') \le S(\pi)$.

Proof. For each $u \in L_0$, the pigeon and functionality axioms derive the row equation

$$h_u = 1 + \sum_{v \in N_H(u)} x_{u,v} = 0$$

in degree at most Δ_L. Indeed, expand the pigeon polynomial and cancel every monomial of degree at least two using a functionality axiom that divides it.

Let Y_u be the set of free variables in row u. The image of the pigeon axiom for $u \in L_0$ is

$$\left(\sum_{y \in Y_u} y \right) \prod_{y \in Y_u} (1 + y) = 0$$

in the Boolean quotient. The pigeon axiom for u_* is unchanged, and all other image axioms have degree at most two. For a quadratic axiom xy, the identity

$$\sigma(xy) + xy = (\sigma(x) + x)y + \sigma(x)(\sigma(y) + y)$$

expresses its difference from the original axiom as a sum of affine multiples of row polynomials. These multiples are derivable in degree at most $\Delta_L + 1$. Thus every axiom of $\mathcal{F}'$ has a degree-$(\Delta_L + 1)$ derivation from $\mathcal{F}$. Prepending these derivations proves (19).

Every Boolean assignment to the free variables extends by (18) to a Boolean assignment to the original variables, so σ preserves semantic inferences. It also takes products of affine forms to products of affine forms. Hence it does not increase the number of distinct non-zero product terms. □

We next construct a space $U_k \subseteq \mathcal{B}^{\leq k}_{n_{\text{free}}}$ such that $U_k \cap D_B(\mathcal{F}') = \{0\}$.

Let E_0 be the set of edges in $E(H) \setminus M$ whose pigeon belongs to L_0, so $|E_0| = n_{\text{free}} - d_*$. For $e = (u, v) \in E_0$, put

$$K(e) = \{p(u), v\},$$
$$W(e) = N_H(p(u)) \cup N_H(v) \subseteq L.$$

We call $A \subseteq E_0$ *compatible* if the sets $W(e)$, for $e \in A$, are pairwise disjoint. For such a set, write $x_A = \prod_{e \in A} x_e$, and define

$$U_k = \operatorname{span}\{x_A : A \subseteq E_0 \text{ compatible}, \ |A| \leq k\}. \tag{20}$$

We will use the following theorem of Mikša and Nordström.

Theorem 6.9 ([54])**.** *Let Γ be a bipartite (s, ε)-boundary expander with left degree at most Δ_L, and let $\mathcal{F}_{\mathsf{FPHP}(\Gamma)}$ be its polynomial encoding over $\mathbb{F}_2$. For every non-negative integer B satisfying*

$$B \leq \frac{\varepsilon s}{2\Delta_L},$$

we have

$$1 \notin D_B(\mathcal{F}_{\mathsf{FPHP}(\Gamma)}),$$

where PC *derivations are interpreted in the Boolean quotient.*

With this theorem, we can prove the following separation lemma.

Lemma 6.10 (Separation for functional PHP)**.** *Let $1 \leq k \leq B$, and suppose that*

$$\Delta_L + 1 \leq B \leq \frac{(\delta - 2)s}{2\Delta_L}.$$

Then

$$U_k \cap D_B(\mathcal{F}') = \{0\}.$$

Proof. Suppose that a non-zero $f \in U_k$ has a degree-B derivation from $\mathcal{F}'$. By Lemma 6.8, it also has such a derivation from $\mathcal{F}$. If $f = 1$, Theorem 6.9 applied to H with $\varepsilon = \delta$ gives a contradiction. Otherwise, choose a monomial x_A of maximum degree in f.

Let I be the set of pigeons in A, and put $K = \bigcup_{e \in A} K(e)$. Since the sets $W(e)$ are pairwise disjoint, the edges of A have distinct pigeons and the pairs $K(e)$ are disjoint. Moreover, every pigeon has at most two neighbours in K.

For each $\beta \in \{0, 1\}^A$, define a restriction $\rho_{A,\beta}$ by setting

$$\begin{aligned} x_{u,v} &= \beta_e, \\ x_{u,p(u)} &= 1 + \beta_e \end{aligned} \qquad (e = (u, v) \in A),$$

setting all other variables in the rows I to zero, and setting $x_{w,z} = 0$ for $w \notin I$ and $z \in K$. Every selected pigeon occupies exactly one of its two reserved holes. Thus all these restrictions leave the same formula $\mathsf{FPHP}(H_A)$, up to zero axioms, where

$$H_A = H - I - K.$$

This graph has $n + 1 - |A|$ pigeons and $n - 2|A|$ holes, so it still gives a pigeonhole contradiction. For every $J \subseteq L \setminus I$ with $|J| \leq s$,

$$\begin{aligned} |\partial_{H_A}(J)| &= |\partial_H(J) \setminus K| \\ &\geq \delta|J| - \sum_{w \in J} |N_H(w) \cap K| \\ &\geq (\delta - 2)|J|. \end{aligned} \tag{21}$$

Hence H_A is an $(s, \delta - 2)$-boundary expander of left degree at most Δ_L.

Apply each restriction $\rho_{A,\beta}$ to the derivation of f. This gives degree-B derivations from the same system $\mathcal{F}_{\mathsf{FPHP}(H_A)}$. Adding their final polynomials derives

$$\sum_{\beta \in \{0,1\}^A} f \restriction \rho_{A,\beta} = 1. \tag{22}$$

To see the equality, the chosen monomial x_A contributes 1, since it evaluates to 1 only when all entries of β are one. Every other monomial of f omits some x_e with $e \in A$, by maximality of the degree of x_A. The only other variable depending on β_e is the matching variable in the same row, which does not occur in f. The restricted monomial is therefore independent of β_e, and its contributions cancel in pairs.

We have obtained a degree-B PC refutation of $\mathsf{FPHP}(H_A)$. By (21), this contradicts Theorem 6.9 with $\Gamma = H_A$ and $\varepsilon = \delta - 2$. □

The degree bounds on H also give a lower bound on the dimension of U_k. Put

$$C = 2\Delta_R \Delta_L (\Delta_R + \Delta_L).$$

Lemma 6.11 (Counting compatible monomials)**.** *Let $k \geq 1$ and suppose that $C(k-1) < n_{\text{free}} - d_*$. Then*

$$\dim(U_k) \geq \frac{1}{k!} \prod_{i=0}^{k-1} (n_{\text{free}} - d_* - Ci) \geq \frac{(n_{\text{free}} - d_* - C(k-1))^k}{k!}. \tag{23}$$

Proof. Fix $e \in E_0$. If $W(e') \cap W(e) \neq \varnothing$ for $e' = (u', v') \in E_0$, then for some $w \in W(e)$ we have $v' \in N_H(w)$ or $p(u') \in N_H(w)$. There are at most $2\Delta_R$ choices for w and at most Δ_L choices for a hole in $N_H(w)$. For each such hole, at most Δ_R edges have it as v'. Also, it is the matching hole of a unique pigeon u', which has at most Δ_L incident edges. Thus at most C edges $e' \in E_0$ have $W(e') \cap W(e) \neq \varnothing$.

After choosing i compatible edges, at least $n_{\text{free}} - d_* - Ci$ choices remain. There are therefore at least $\prod_{i=0}^{k-1}(n_{\text{free}} - d_* - Ci)$ ordered choices of k compatible edges. Dividing by $k!$ counts distinct monomials, which are linearly independent in the free variables. The second inequality follows by bounding each factor below by $n_{\text{free}} - d_* - C(k-1)$. □

We can now apply the rank bound from Theorem 4.4.

Theorem 6.12 (Functional graph PHP)**.** *Fix positive integers $\Delta_{L,0}, \Delta_{R,0}$ and constants $\delta_0 > 3$ and $\gamma > 0$. Suppose that G has $n+1$ pigeons, n holes, left degree at most $\Delta_{L,0}$, right degree at most $\Delta_{R,0}$, and is a $(\gamma n, \delta_0)$-boundary expander. For any matching M covering all holes, every semantic $\Sigma\Pi\Sigma$-PC refutation π of $\mathsf{FPHP}(G + M)$ satisfies*

$$S(\pi) \geq 2^{\Omega(n)}.$$

The implicit constant depends only on the fixed graph parameters. The same lower bound holds for total $\Sigma\Pi\Sigma$-size.

Proof. Use the notation above, with $s = \gamma n$. Since $\delta > 2$, every pigeon of H has degree at least three. Thus $n_{\text{free}} = \Theta(n)$, while $d_* \leq \Delta_L$ and C are constants. Choose $k = \lfloor \eta n \rfloor$ for a sufficiently small constant $\eta > 0$, and put

$$r = (C+1)k + d_*,$$
$$B = 3(k + r).$$

For sufficiently large n, the choice of η ensures that $r \geq \Delta_L$, $B \leq n_{\text{free}}$, and

$$B \leq \frac{(\delta - 2)\gamma n}{2\Delta_L}.$$

In particular, $B = 3(k + \max\{\Delta_L, r\})$ and $B \geq \Delta_L + 1$.

By Lemma 6.8, applying σ gives a refutation π' of $\mathcal{F}'$ with $S(\pi') \le S(\pi)$. Let $\mathcal{R}$ be the naïve degree-B pseudo-reduction operator for $\mathcal{F}'$ from Proposition 4.3. Its kernel is $D_B(\mathcal{F}')$, so Lemma 6.10 gives

$$\operatorname{rank}(\mathcal{R}_{\le k}) \ge \dim(U_k).$$

The axioms of $\mathcal{F}'$ have degree at most Δ_L. Hence Theorem 4.4 and Lemma 6.11 give

$$\begin{aligned} S(\pi) \ge S(\pi') &\ge \frac{\dim(U_k)}{2\binom{n_{\text{free}}-r}{k}} \\ &\ge \frac{1}{2}\left(\frac{n_{\text{free}} - d_* - C(k-1)}{n_{\text{free}} - r}\right)^k \\ &= \frac{1}{2}\left(1 + \frac{k+C}{n_{\text{free}} - r}\right)^k \\ &\ge \frac{1}{2}\left(1 + \frac{k}{n_{\text{free}}}\right)^k \\ &\ge \frac{1}{2}\exp\left(\frac{k^2}{2n_{\text{free}}}\right) = 2^{\Omega(n)}. \end{aligned}$$

Here $B \le n_{\text{free}}$ ensures that all denominators are positive and that $C(k-1) < n_{\text{free}} - d_*$. We used $\binom{n_{\text{free}}-r}{k} \le (n_{\text{free}} - r)^k/k!$ and $\ln(1+x) \ge x/2$ for $0 \le x \le 1$. The same bound holds for total circuit size, which is at least $S(\pi)$. □

Corollary 6.13 (Complete functional PHP)**.** *Every semantic* ΣΠΣ-PC *refutation* π *of the complete principle* FPHP_n^{n+1} *satisfies*

$$S(\pi) \ge 2^{\Omega(n)}.$$

Proof. Take G from Lemma 6.7 and add any matching M covering all holes. Setting $x_{u,v} = 0$ for every edge outside $G + M$ restricts FPHP_n^{n+1} to $\mathsf{FPHP}(G + M)$. Restrictions preserve semantic inferences and do not increase $S(\pi)$, so Theorem 6.12 gives the lower bound. □

We use an affine substitution to transfer the exponential lower bound for $\mathsf{FPHP}(G+M)$ to semantic ΣΠΣ-PC refutations of the bit pigeonhole principle BPHP_n^{n+1}, defined in Section 3.

Corollary 6.14 (Bit PHP)**.** *Let* $n = 2^\ell$*. Every semantic* ΣΠΣ-PC *refutation* π *of* BPHP_n^{n+1} *satisfies*

$$S(\pi) \ge 2^{\Omega(n)}.$$

The same lower bound holds for total ΣΠΣ*-size.*

Proof. Take G from Lemma 6.7 and add a matching M covering all holes. Put $H = G+M$ and identify its holes with $\{0,1\}^\ell$. For the bit variables $y_{u,t}$ of BPHP_n^{n+1}, define

$$\tau(y_{u,t}) = \sum_{\substack{v \in N_H(u) \\ v_t = 1}} x_{u,v}.$$

On every assignment satisfying the pigeon and functionality axioms for row u, these affine forms give the binary address of its unique occupied hole. The substitution τ preserves semantic inferences and does not increase the number of distinct non-zero product terms.

We show that each substituted collision axiom has a semantic ΣΠΣ-PC derivation from $\mathsf{FPHP}(H)$ using a constant number of product terms. Fix pigeons u, u', and let $g_1, \dots, g_t$ be their pigeon and functionality axioms together with the hole axioms involving both rows. There are $O(\Delta_L^2) = O(1)$ such axioms on at most $2\Delta_L$ variables. Their conjunction implies every substituted collision axiom for u, u', since the two occupied holes must differ.

Derive

$$Q_j = 1 + \prod_{i=1}^{j}(1 + g_i),$$

starting with $Q_1 = g_1$ and deriving Q_j from Q_{j-1} and g_j for $j \geq 2$. Each step is a sound semantic inference, since $Q_j = 0$ expresses that $g_1, \ldots, g_j$ all vanish. Represent each Q_j by its multilinear expansion. Since it depends on at most $2\Delta_L$ variables, this expansion has at most $2^{2\Delta_L}$ monomials and gives a $\Sigma\Pi\Sigma$ representation of constant size. From Q_t, derive the substituted collision axiom by a semantic inference.

There are $O(n^3)$ collision axioms. Derive their images as above and then apply τ to the refutation π. The resulting refutation π_H of $\mathsf{FPHP}(H)$ satisfies

$$S(\pi_H) \leq S(\pi) + O(n^3).$$

The claim now follows from Theorem 6.12. □

For $n+1$ pigeons, our bound matches Pang's $2^{\Omega(n)}$ bound for $\mathsf{Res}(\oplus)$ [55] and extends it to semantic $\Sigma\Pi\Sigma$-PC.

6.3 Semantic $\Sigma\Pi\Sigma\Pi^{[2]}$-PC is strictly stronger than $\Sigma\Pi\Sigma$-PC

We now combine the preceding lower bounds with polynomial-size refutations in semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC. We first observe that semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC efficiently refutes every unsatisfiable system of degree at most w over $\mathbb{F}_2$. For $w = 3$, this gives an exponential separation on random 3-CNFs. We then give a polynomial-size refutation of the functional pigeonhole principle in semantic $\Sigma\Pi\Sigma\Pi^{[2]}$-PC, yielding an exponential separation already for $w = 2$.

For a polynomial $f \in \mathcal{B}_n(\mathbb{F}_2)$, let $\mathrm{size}_{\mathrm{sp}}(f)$ denote the size of its natural sparse $\Sigma\Pi$ representation of its unique multilinear representative, and put

$$\mathrm{size}_{\mathrm{sp}}(\mathcal{F}) = \sum_{f \in \mathcal{F}} \mathrm{size}_{\mathrm{sp}}(f).$$

Proposition 6.15 (Bounded-degree systems)**.** *Let*

$$\mathcal{F} = \{f_1, \ldots, f_t\} \subseteq \mathcal{B}_n(\mathbb{F}_2)$$

be a collection of polynomials of degree at most w such that the equations

$$f_1 = 0, \ldots, f_t = 0$$

have no common solution in $\{0,1\}^n$. Then $\mathcal{F}$ has a semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutation with at most $2t-1$ lines and of $\Sigma\Pi\Sigma\Pi^{[w]}$-PC-size

$$O\big((t + \mathrm{size}_{\mathrm{sp}}(\mathcal{F}))\log(t+1)\big).$$

In particular, for constant w this size is polynomial in n and t.

Proof. For every nonempty $A \subseteq [t]$, define

$$Q_A = 1 + \prod_{i \in A}(1 + f_i). \tag{24}$$

Since the ground field is $\mathbb{F}_2$, for every Boolean assignment a we have

$$Q_A(a) = 0 \quad \Longleftrightarrow \quad f_i(a) = 0 \text{ for every } i \in A. \tag{25}$$

Moreover, Q_A is computed by a $\Sigma\Pi\Sigma\Pi^{[w]}$ circuit: each f_i is a sum of monomials of degree at most w, each factor $1 + f_i$ is therefore a $\Sigma\Pi^{[w]}$ circuit, and (24) is one product of such factors plus the constant 1.

Arrange the indices in $[t]$ as the leaves of a balanced binary tree. The leaf corresponding to i is labelled by $Q_{\{i\}} = f_i$, which is an axiom. If two children of an internal node are labelled by disjoint sets A and B, derive $Q_{A \cup B}$ from Q_A and Q_B. This is a sound semantic inference: if both premises

vanish at a Boolean assignment, then all polynomials with indices in $A \cup B$ vanish there, and hence so does $Q_{A\cup B}$.

At the root we obtain $Q_{[t]}$. Since the polynomials in $\mathcal{F}$ have no common Boolean zero, (25) implies that $Q_{[t]}(a) = 1$ for every $a \in \{0,1\}^n$. Thus $Q_{[t]} = 1$ in $\mathcal{B}_n(\mathbb{F}_2)$, and the root line is a refutation.

The tree has at most $2t-1$ nodes. The circuit at a node labelled by A has size

$$O\left(|A| + \sum_{i\in A} \mathrm{size}_{\mathrm{sp}}(f_i)\right).$$

At each level the sets labelling the nodes are pairwise disjoint, and the tree has height $O(\log(t+1))$. Summing over all levels gives the claimed size bound. □

Corollary 6.16 (Upper bound for bounded-width CNFs)**.** *Let $w \geq 1$ and let F be an unsatisfiable CNF with m clauses, each of width at most w. Then F has a semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutation over $\mathbb{F}_2$ of total circuit size*

$$O(mw2^w \log(m+1)).$$

In particular, for constant w the size is $O(m\log(m+1))$.

Proof. Every clause polynomial has degree at most w, and its multilinear expansion has at most 2^w monomials, each of degree at most w. Thus its sparse representation has size $O(w2^w)$. Apply Proposition 6.15 to the clause polynomials. □

Corollary 6.17 (Separation on random 3-CNFs)**.** *For every fixed $\alpha > \ln 2/\ln(8/7)$, with probability $1-o(1)$ the random formula $F_{n,\alpha}$ is unsatisfiable, has a semantic $\Sigma\Pi\Sigma\Pi^{[3]}$-PC refutation over $\mathbb{F}_2$ of total circuit size $O_\alpha(n\log n)$, and every semantic $\Sigma\Pi\Sigma$-PC refutation π of $F_{n,\alpha}$ satisfies*

$$S(\pi) \geq 2^{\Omega_\alpha(n)}.$$

Proof. The unsatisfiability and lower bound follow from Theorem 6.5. Since $F_{n,\alpha}$ has $\lfloor \alpha n \rfloor$ clauses of width three, Corollary 6.16 gives the upper bound. □

We next apply Proposition 6.15 to the functional pigeonhole principle from Definition 6.6.

Theorem 6.18 (Functional pigeonhole principle)**.** *The principle FPHP_n^{n+1} has a semantic $\Sigma\Pi\Sigma\Pi^{[2]}$-PC refutation over $\mathbb{F}_2$ of size*

$$O(n^3 \log(n+1)).$$

Consequently, the same upper bound holds in semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC for every $w \geq 2$.

Proof. We first derive, for every pigeon $u \in [n+1]$, the row-sum polynomial

$$R_u := 1 + \sum_{v\in[n]} x_{u,v}. \tag{26}$$

Apply the construction from the proof of Proposition 6.15 to the functionality axioms belonging to u. It yields the polynomial

$$A_u := 1 + \prod_{\{v,v'\}\in\binom{[n]}{2}} (1 + x_{u,v}x_{u,v'}).$$

Its zero set consists precisely of the assignments satisfying all functionality axioms for pigeon u. The pigeon axiom P_u is itself representable by a $\Sigma\Pi\Sigma\Pi^{[2]}$-PC circuit: each affine factor $1 + x_{u,v}$ can be placed at the third layer. Hence we may make the semantic inference

$$\frac{A_u \qquad P_u}{R_u}.$$

Indeed, $A_u = 0$ says that at most one variable in row u is equal to one, whereas $P_u = 0$ says that at least one of them is equal to one. Thus exactly one variable in the row is one, which implies $R_u = 0$.

Consider now the system

$$\mathcal{Q}_n = \{R_u : u \in [n+1]\} \cup \{F_{u;v,v'} : u \in [n+1],\ v < v'\} \cup \{H_{u,u';v} : u < u',\ v \in [n]\}.$$

All polynomials in $\mathcal{Q}_n$ have degree at most two. They have no common Boolean zero: the row-sum and functionality equations make every pigeon choose exactly one hole, while the hole equations make these choices injective, which is impossible for $n+1$ pigeons and n holes. Therefore, Proposition 6.15 applied to the already available lines in $\mathcal{Q}_n$ yields a semantic $\Sigma\Pi\Sigma\Pi^{[2]}$-PC refutation.

The system $\mathcal{Q}_n$ contains

$$(n+1)\binom{n}{2} + n\binom{n+1}{2} + (n+1) = O(n^3)$$

polynomials, with total sparse size $\mathrm{size}_{\mathrm{sp}}(\mathcal{Q}_n) = O(n^3)$. Hence the balanced-tree construction has size $O(n^3 \log(n+1))$. The preliminary derivations of all polynomials A_u and R_u have total size $O(n^3 \log(n+1))$ as well. □

Remark 6.19. *The preservation of the bottom fan-in w in Proposition 6.15 uses the fact that the ground field is $\mathbb{F}_2$. Over a finite field $\mathbb{F}_q$, the analogous zero-test construction replaces f_i by f_i^{q-1} and gives bottom fan-in at most $w(q-1)$ after expansion.*

Corollary 6.20 (Separation on the functional pigeonhole principle)**.** *The functional pigeonhole principle FPHP_n^{n+1} has a semantic $\Sigma\Pi\Sigma\Pi^{[2]}$-PC refutation over $\mathbb{F}_2$ of total circuit size $O(n^3 \log(n+1))$, whereas every semantic $\Sigma\Pi\Sigma$-PC refutation π satisfies $S(\pi) \geq 2^{\Omega(n)}$. Consequently, semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC is strictly stronger than semantic $\Sigma\Pi\Sigma$-PC for every $w \geq 2$.*

Proof. Combine Corollary 6.13 with Theorem 6.18. Every $\Sigma\Pi\Sigma$-PC line is also a $\Sigma\Pi\Sigma\Pi^{[1]}$-PC line, and every $\Sigma\Pi\Sigma\Pi^{[2]}$-PC line is a $\Sigma\Pi\Sigma\Pi^{[w]}$-PC line for $w \geq 2$, so the same separation holds for every $w \geq 2$. □

7 Lower Bounds for Semantic $\Sigma\Pi\Sigma\Pi$-PC

This section proves lower bounds for semantic $\Sigma\Pi\Sigma\Pi$-PC refutations. We first consider refutations with bounded bottom degree over $\mathbb{F}_2$, allowing any semantically sound inference from at most two premises, as defined in the preliminaries. As before, $S(\pi)$ counts distinct non-zero top product terms, and every lower bound on $S(\pi)$ also holds for total circuit size.

In Section 7.1, we prove lower bounds for semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutations of random logarithmic-width CNFs. In Section 7.2, we construct sparse boundary-expanding graph instances of the unary pigeonhole principle and prove corresponding lower bounds for semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC. Both arguments use the rank estimates for the Mikša–Nordström operator from Section 5.

We then remove the bound on the bottom degree in two ways. In Section 7.3, a selector lift transfers lower bounds for bounded bottom degree to circuit-size lower bounds with unrestricted bottom degree. The lifting theorem and its consequences apply to any base CNF over any field; we apply them to the preceding random formulas. In Section 7.4, we sparsify a putative small refutation of the complete pigeonhole principle, reducing it to one of the graph instances with bounded bottom degree. This also gives a superpolynomial lower bound for Res(log) refutations of the weak pigeonhole principle.

7.1 Random $O(\log n)$-CNFs in semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC

Put $\ell = \lceil \log_2 n \rceil$ and $m = n2^\ell$, and let F_n consist of m independently and uniformly sampled clauses of width exactly ℓ, each on distinct variables. Equivalently, each clause is sampled from the $2^\ell\binom{n}{\ell}$ possible clauses, with repetition. Thus F_n has $\Theta(n^2)$ clauses and $O(n^2 \log n)$ literal occurrences. We again identify it with its standard polynomial encoding when discussing algebraic proofs.

Lemma 7.1 (Unsatisfiability and expansion)**.** *There is an absolute constant $\eta > 0$ such that, with probability $1 - o(1)$, F_n is unsatisfiable and its incidence graph satisfies*

$$|N_{F_n}(J)| \geq \frac{3\ell}{4}|J|,$$
$$|\partial_{F_n}(J)| \geq \frac{\ell}{2}|J|$$

whenever $1 \leq |J| \leq \lfloor \eta n/\ell \rfloor$.

Both assertions follow from standard first-moment estimates. For completeness, we give a proof in Appendix C.

Theorem 7.2 (Random logarithmic-width CNFs)**.** *There are absolute constants $a > 0$, $C > 1$, and $c > 0$ such that, with probability $1 - o(1)$, F_n is unsatisfiable and simultaneously for every integer $1 \leq w \leq c \log_2 n$, every semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutation π over $\mathbb{F}_2$ satisfies*

$$S(\pi) \geq \frac{1}{4} \exp\left(\frac{an}{wC^w} \right). \tag{27}$$

The same lower bound holds for total circuit size. The constants are independent of both n and w.

Proof. Work on the event in Lemma 7.1, and let $s = \lfloor \eta n/\ell \rfloor$. Choose a sufficiently large absolute constant C and put

$$k = \left\lfloor \frac{4n}{C} \right\rfloor.$$

Since $\ell s = \eta n - O(\ell)$, for sufficiently large n we have

$$\ell \leq k, \qquad 6k \leq \min\{n, \ell s/4 - 1\}.$$

Thus Lemma 5.3 supplies an expansion pseudo-reduction operator $\mathcal{R}$ defined up to degree $6k$.

We apply Theorem 5.4 with singleton right clusters, $D = 1$, boundary-expansion factor $\ell/2$, and $\kappa = m$, which bounds the degree of the graph on the m clauses. Since $m = n2^\ell$,

$$\varepsilon = m(k/n)^{\ell/2} \leq n \left(\frac{4}{\sqrt{C}} \right)^\ell = o(1)$$

for sufficiently large C, because $\ell \geq \log_2 n$. In particular, $4\varepsilon < 1$, and the theorem gives

$$\operatorname{rank}(\mathcal{R}_{\leq k}) \geq \left(1 - \frac{\varepsilon}{1 - 4\varepsilon}\right) \binom{n}{k} \geq \frac{1}{2} \binom{n}{k}. \tag{28}$$

Now fix an integer $1 \leq w \leq \log_2 n$ and put

$$r = \left\lfloor \frac{k}{w} \right\rfloor.$$

The parameters satisfy $2w \leq k$ and

$$3(k + \max\{\ell, rw\}) \leq 6k.$$

We can therefore apply Corollary 4.6, using the restriction of $\mathcal{R}$ to the required degree. By (28), every semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutation π satisfies

$$S(\pi) \geq \frac{1}{4} \exp\left(\frac{r}{4} \left(\frac{k}{2n} \right)^w \right).$$

For sufficiently large n,

$$\frac{k}{2n} \geq \frac{1}{C}, \qquad r \geq \frac{k}{2w} \geq \frac{n}{Cw}.$$

Consequently,

$$S(\pi) \geq \frac{1}{4} \exp\left(\frac{n}{4Cw} C^{-w} \right).$$

This is the claimed bound with $a = 1/(4C)$. Choose $c = 1/(2 \log_2 C)$.

All estimates hold on the same expansion event, simultaneously for every $1 \leq w \leq c \log_2 n$. Unsatisfiability follows from Lemma 7.1, and the total circuit size is at least $S(\pi)$. □

Corollary 7.3. *With probability* $1 - o(1)$*, simultaneously for all* $1 \le w \le c \log_2 n$*, every semantic* $\Sigma\Pi\Sigma\Pi^{[w]}$-PC *refutation of* F_n *has*

$$S(\pi) \ge 2^{\Omega(\sqrt{n}/\log n)}.$$

If $w = o(\log n)$*, the more precise bound is* $S(\pi) \ge 2^{n^{1-o(1)}}$.

Proof. Our choice of c gives $C^w \le \sqrt{n}$. Substitute this in (27). If $w = o(\log n)$, then $wC^w = n^{o(1)}$, yielding the second assertion. □

7.2 Graph PHP in semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC

This subsection proves lower bounds for the unary pigeonhole principle in semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC.

Let $G = (L \sqcup R, E)$ be a bipartite graph, where L is a set of pigeons and R is a set of holes. For every edge $(u, v) \in E$, the variable $x_{u,v}$ indicates that pigeon u is mapped to hole v. The graph pigeonhole principle $\mathsf{PHP}(G)$ consists of the axioms

$$P_u = \prod_{v \in N_G(u)} (1 - x_{u,v}) = 0 \qquad u \in L,$$

$$x_{u,v} x_{u',v} = 0 \qquad u, u' \in L,\ u \ne u',\ v \in N_G(u) \cap N_G(u').$$

Thus every pigeon must occupy an adjacent hole, while no hole may contain two pigeons. We do not impose functionality axioms. The standard unary pigeonhole principle PHP^m_n is $\mathsf{PHP}(K_{m,n})$.

We use the expansion framework from Section 5, taking the pigeon axioms as the left clusters and the variables associated with each hole as a right cluster. The hole axioms form $\mathcal{F}_0$. As observed there, this clustered representation is respectful: if v is a boundary hole with unique neighbouring pigeon u, setting $x_{u,v} = 1$ and all other variables associated with v to zero satisfies P_u and preserves the hole axioms.

We first construct the expanding graphs used in both lower bounds. Recall that a bipartite graph $G = (L \sqcup R, E)$ is an (s, c)-boundary expander if, for every set $I \subseteq L$ of pigeons with $|I| \le s$, its boundary $\partial_G(I)$ satisfies $|\partial_G(I)| \ge c|I|$.

Lemma 7.4 (Boundary expanders for PHP)**.** *For all sufficiently large* n*, put*

$$\Delta = \lceil 24 \ln(2n) \rceil$$

and

$$s = \left\lfloor \frac{n}{4\Delta} \right\rfloor.$$

Let G *be a random bipartite graph with* $2n$ *pigeons and* n *holes in which every pigeon–hole pair is an edge independently with probability* Δ/n*. Asymptotically almost surely,* G *is an* $(s, \Delta/4)$*-boundary expander, every pigeon has degree at most* 2Δ*, and every hole has degree at most* 4Δ*.*

The proof is given in Appendix D.

To apply Theorem 4.5, we need to partition the variables into n equal-sized classes and include in $\mathcal{F}_0$ all products of two distinct variables from the same class. The edge variables incident with a hole form a natural class, but the right degrees need not be equal. Fix a graph G satisfying the conclusions of Lemma 7.4, put $D = 4\Delta$, and pad the class associated with every hole v with auxiliary variables to obtain a class C_v of size D. Let

$$\mathcal{F}_0 = \bigcup_{v \in R} \{xy : x, y \in C_v,\ x \ne y\}$$

and consider the polynomial system

$$\mathsf{PHP}'(G) = \mathcal{F}_0 \cup \{P_u : u \in L\}. \tag{29}$$

The original hole axioms belong to $\mathcal{F}_0$, and hence $\mathsf{PHP}(G) \subseteq \mathsf{PHP}'(G)$. The auxiliary variables occur only in the additional hole axioms, and setting them to zero recovers $\mathsf{PHP}(G)$. Thus, $\mathsf{PHP}'(G)$ is a

trivial extension of $\mathsf{PHP}(G)$, introduced only to make all hole classes C_v have the same size D, as required by Theorem 4.5. With this choice of $\mathcal{F}_0$, the algebra $\mathcal{B}_{\mathcal{F}_0}$ has a basis consisting of monomials that contain at most one variable associated with each hole.

The clustered representation with left clusters $\{\{P_u\} : u \in L\}$ and right variable clusters $\{C_v : v \in R\}$ is respectful, as observed above. We can therefore apply the operator construction from Section 5 in $\mathcal{B}_{\mathcal{F}_0}$.

Theorem 7.5 (Bounded-bottom-degree graph PHP)**.** *There are absolute constants $c_0, C_0 > 0$ such that the following holds for all sufficiently large n. Let G satisfy the conclusions of Lemma 7.4. Simultaneously for every integer*

$$1 \le w \le \frac{\ln n}{\ln \ln n},$$

every semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutation π of $\mathsf{PHP}(G)$ satisfies

$$S(\pi) \ge \frac{1}{4} \exp\left(\frac{c_0 n}{w\Delta(C_0\Delta^2)^w} \right).$$

The same lower bound holds for the following more general form of a semantic refutation. Semantic inferences need only be sound on the hole domain

$$\Omega_{\mathcal{F}_0} = \{\alpha \in \{0,1\}^{nD} : f(\alpha) = 0 \text{ for every } f \in \mathcal{F}_0\},$$

and every initial line need only agree on $\Omega_{\mathcal{F}_0}$ either with an axiom of $\mathsf{PHP}'(G)$ or with the zero polynomial. Equivalently, in $\mathcal{B}_{\mathcal{F}_0}$, every initial line is equal either to the image of an axiom or to zero.

Proof. Put

$$k = \left\lfloor \frac{s}{512} \right\rfloor, \qquad B_0 = \left\lfloor \frac{s}{32} \right\rfloor.$$

For sufficiently large n,

$$2\Delta \le B_0 \le \min\{n, (\Delta/4)s/2 - 1\}.$$

Thus Definition 5.2 and Lemma 5.3 give a degree-B_0 expansion operator $\mathcal{R}$ for $\mathsf{PHP}'(G)$ in $\mathcal{B}_{\mathcal{F}_0}$.

We first estimate its rank using Theorem 5.4. There are $2n$ left vertices and n right clusters of size D; every monomial choosing at most one variable from each cluster is standard. The boundary-expansion factor is $\Delta/4$, and we may take $\kappa = 2\Delta(D-1) = O(\Delta^2)$. Since $s = \Theta(n/\Delta)$,

$$p = \frac{k}{n} = \Theta(1/\Delta), \qquad p^{\Delta/4} = n^{-\Theta(\ln \ln n)}.$$

Consequently, both $4\kappa p^{\Delta/4}$ and $2np^{\Delta/4}$ tend to zero, and the general rank estimate gives

$$\operatorname{rank}(\mathcal{R}_{\le k}) \ge \left(1 - \frac{2np^{\Delta/4}}{1 - 4\kappa p^{\Delta/4}}\right) \binom{n}{k} D^k \ge \frac{1}{2}\binom{n}{k} D^k \tag{30}$$

for all sufficiently large n.

A refutation of $\mathsf{PHP}(G)$ remains a semantic refutation on the domain of $\mathcal{F}_0$, with the auxiliary variables and additional hole axioms unused. Passing to this quotient can only decrease $S(\pi)$. It therefore suffices to prove the bound for the more general refutations in the statement.

Fix w in the stated range and put

$$r = \left\lfloor \frac{B_0}{6w} \right\rfloor.$$

The maximum axiom degree is at most 2Δ, and, for all sufficiently large n, $2w \le k$ and

$$3(k + \max\{2\Delta, rw\}) \le B_0 \le n.$$

Apply Theorem 4.5 to the restriction of $\mathcal{R}$ to the required degree, with n groups of size D. By (30),

$$\begin{aligned} S(\pi) &\ge \frac{\operatorname{rank}(\mathcal{R}_{\le k})}{2\binom{n}{k}D^k} \exp\left(\frac{r}{4}\left(\frac{k}{2nD}\right)^w\right) \\ &\ge \frac{1}{4}\exp\left(\frac{r}{4}\left(\frac{k}{2nD}\right)^w\right). \end{aligned}$$

Since $s = \Theta(n/\Delta)$ and $D = 4\Delta$, there are absolute positive constants c_1, c_2 such that

$$r \geq \frac{c_1 n}{w\Delta}, \qquad \frac{k}{2nD} \geq \frac{1}{c_2\Delta^2}.$$

Absorbing these constants into c_0 and C_0 proves the result. □

7.3 Lifting to semantic ΣΠΣΠ-PC

The first two subsections give lower bounds with bounded bottom degree. We now give a lifting theorem for an arbitrary base CNF. We start with a formula that is hard for ΣΠΣΠ-PC with bounded bottom fan-in and construct a lifted formula that reduces to the original formula under a large family of partial assignments, which we call *projections*. We show that if the lifted formula has a small refutation in unrestricted semantic ΣΠΣΠ-PC, then a random projection transforms this refutation, with positive probability, into a refutation of the original formula in which all bottom product gates have bounded fan-in. We can then apply the lower bound for the bounded-bottom-fan-in system.

For each variable of a base CNF F, the lifted formula has a block with q cells of each of two colours. Selector and exclusion clauses require exactly one chosen cell in every block; its colour encodes the base variable's Boolean value. Thus $\mathrm{Lift}_q(F)$ is unsatisfiable whenever F is. If F has polynomial size and q is polynomial in its number of variables, the lifted CNF also has polynomial size.

A random projection chooses one cell of each colour in every block, maps these two cells to a base literal and its complement, and sets all other cells to zero. The lifted input equations project to the base input equations or to $0 = 0$. Each Boolean assignment to the base variables induces a Boolean assignment to the lifted variables, so semantic inferences remain sound and retain their original arity.

The degree compression comes from monomial survival: a bottom monomial using s blocks survives with probability at most q^{-s}. If a source proof contains R distinct bottom monomials and $R < q^{d+1}$, a union bound supplies one projection under which every bottom factor has degree at most d. We keep all outer products intact. Since the same ring homomorphism acts on every proof line, the number of distinct non-zero top product terms cannot increase, even when source products are identified as Boolean functions.

This gives a transfer involving two resources. If $B_F(d)$ is a lower bound on distinct top product terms in semantic $\Sigma\Pi\Sigma\Pi^{[d]}$-PC refutations of F, then any unrestricted refutation of $\mathrm{Lift}_q(F)$ must have at least q^{d+1} distinct bottom monomials or at least $B_F(d)$ distinct top product terms. Total circuit size bounds both quantities. For our random formulas, choosing $q = n$ and $d = \Theta(\log n)$ yields the $n^{\Omega(\log n)}$ circuit-size lower bound in Theorem 7.12.

The lifting itself works over every field. Its application to the random-formula lower bound is over $\mathbb{F}_2$.

7.3.1 Representations and the size measure

For Boolean variables Z and a field $\mathbb{F}$, write

$$\mathcal{B}_Z(\mathbb{F}) = \mathbb{F}[Z]/\langle z^2 - z : z \in Z\rangle.$$

A represented line has the form

$$P = \sum_{a=1}^{t} c_a \prod_{b=1}^{r_a} f_{a,b}, \tag{31}$$

where $c_a \in \mathbb{F}$ and each bottom polynomial $f_{a,b}$ is written as a sum of monomials. Requiring $\deg f_{a,b} \leq d$ gives $\Sigma\Pi\Sigma\Pi^{[d]}$-PC; degrees are measured after multilinearisation. Empty products and constant bottom factors are allowed.

We use the semantic convention from the preliminaries: an inference from one or two earlier equations is allowed if it is sound on the Boolean cube. Initial equations are input clause equations or Boolean tautologies, including $0 = 0$, and a refutation ends with $1 = 0$. The bound on inference arity is retained throughout the projection.

Let $M_{\mathrm{bot}}(\pi)$ be the maximum of 1 and the number of distinct monomials occurring in the multilinear bottom polynomials of π. Continue to use $S(\pi)$ for the number of distinct non-zero top product terms, identified in the Boolean quotient. If $V(\pi)$ is the set of proof vertices and C_v is the displayed circuit at v, write

$$|\pi|_4 = \sum_{v \in V(\pi)} |C_v| \tag{32}$$

for the total $\Sigma\Pi\Sigma\Pi$ circuit size of the proof. Both measures satisfy

$$\begin{aligned} M_{\mathrm{bot}}(\pi) &\le |\pi|_4, \\ S(\pi) &\le |\pi|_4. \end{aligned} \tag{33}$$

Indeed, a bottom monomial is computed by a bottom product gate, or is a variable or constant input, while a top product term is computed by a second-layer product gate. Shared gates do not affect these bounds, since the two measures count distinct polynomials rather than occurrences.

7.3.2 The selector formula and its projections

Let F be a CNF on $y_1, \dots, y_n$, and let $q \ge 2$ be an integer. Introduce variables

$$X_{i,b,a} \qquad (i \in [n],\ b \in \{0,1\},\ a \in [q]).$$

The formula $\mathrm{Lift}_q(F)$ consists of:

1. the selector clause $\bigvee_{b \in \{0,1\},\, a \in [q]} X_{i,b,a}$ for each i;
2. the exclusion clause $\neg X_{i,b,a} \vee \neg X_{i,b',a'}$ for every pair of distinct cells $(b,a), (b',a')$ in the same block i;
3. for each $C \in F$, the clause obtained by replacing y_i with $\bigvee_{a \in [q]} X_{i,1,a}$ and $\neg y_i$ with $\bigvee_{a \in [q]} X_{i,0,a}$.

The first two types require exactly one chosen cell in each block; its colour b records the value of y_i. Thus the lift is unsatisfiable exactly when F is unsatisfiable. If F has polynomial size and q is polynomially bounded, then $\mathrm{Lift}_q(F)$ also has polynomial size.

If F has m clauses and Λ literal occurrences, the lift has

$$\begin{aligned} &2nq && \text{variables,} \\ &m + n + n\binom{2q}{2} && \text{clauses,} \\ &q\Lambda + 2nq + 2n\binom{2q}{2} && \text{literal occurrences.} \end{aligned} \tag{34}$$

Let

$$A = (A_{i,b})_{i \in [n],\, b \in \{0,1\}} \in [q]^{[n] \times \{0,1\}}$$

be a random variable whose coordinates $A_{i,b}$ are chosen independently and uniformly from $[q]$. For every realisation $a = (a_{i,b})$ of A, define a map on the variables $X_{i,b,j}$ by

$$X_{i,0,j} \longmapsto \begin{cases} 1 - y_i, & j = a_{i,0}, \\ 0, & j \ne a_{i,0}, \end{cases} \qquad X_{i,1,j} \longmapsto \begin{cases} y_i, & j = a_{i,1}, \\ 0, & j \ne a_{i,1}. \end{cases} \tag{35}$$

This map on the generators extends uniquely to an $\mathbb{F}$-algebra homomorphism

$$\widetilde{\rho}_a : \mathbb{F}[X] \longrightarrow \mathcal{B}_Y(\mathbb{F})$$

by substitution: for an arbitrary polynomial P, replace every variable $X_{i,b,j}$ by its image in (35) and perform the resulting additions and multiplications in $\mathcal{B}_Y(\mathbb{F})$. Since all the images in (35) are Boolean elements, we have

$$\widetilde{\rho}_a(X_{i,b,j}^2 - X_{i,b,j}) = 0$$

for every i, b, j. Hence $\widetilde{\rho}_a$ annihilates the Boolean ideal and factors uniquely through the quotient, giving a ring homomorphism

$$\rho_a : \mathcal{B}_X(\mathbb{F}) \longrightarrow \mathcal{B}_Y(\mathbb{F}).$$

We write ρ_A for the random homomorphism obtained by sampling A, and call it a *random projection*. Thus the only randomness in ρ_A comes from the choice of the coordinates $A_{i,b}$; once A is fixed, ρ_A is a deterministic ring homomorphism.

For every Boolean assignment y, the substitution corresponding to A produces a Boolean assignment $X = X_A(y)$ with exactly one selected cell in each block. Consequently,

$$\rho_A(P)(y) = P(X_A(y)) \tag{36}$$

for every $P \in \mathcal{B}_X(\mathbb{F})$.

Lemma 7.6 (Projection of the input)**.** *For every A, the polynomial of the lifted clause corresponding to $C \in F$ projects to $[\![C]\!]$. The polynomials of all selector and exclusion clauses project to zero.*

Proof. The product $\prod_a(1 - X_{i,1,a})$ projects to $1 - y_i$, and $\prod_a(1 - X_{i,0,a})$ projects to y_i. These are exactly the factors in $[\![C]\!]$. A selector polynomial projects to $y_i(1-y_i) = 0$. For an exclusion polynomial, two cells of the same colour cannot both survive, while surviving cells of opposite colours have product $y_i(1 - y_i) = 0$. □

7.3.3 Compression and proof lifting

Call a multilinear monomial *block-simple* if it uses at most one cell from each block i.

Lemma 7.7 (Monomial survival)**.** *A monomial that is not block-simple is killed by every ρ_A. A block-simple monomial using s blocks survives with probability q^{-s}; whenever it survives, its image has degree exactly s.*

Proof. Two distinct cells from the same block have zero product after projection, as in Lemma 7.6. Otherwise write $m = \prod_{j=1}^{s} X_{i_j,b_j,a_j}$ with distinct i_j. Survival requires the independent events $A_{i_j,b_j} = a_j$. The surviving image is a product of literals on s distinct variables, whose degree-s coefficient is 1 or -1. □

Lemma 7.8 (Simultaneous compression of bottom polynomials)**.** *Let $\mathcal{P}$ be a finite family of bottom polynomials with R distinct monomials in their multilinear representatives. For every integer $d \geq 0$,*

$$\Pr_A[\exists f \in \mathcal{P} : \deg \rho_A(f) > d] \leq Rq^{-(d+1)}. \tag{37}$$

In particular, if $R < q^{d+1}$, a common projection sends every bottom polynomial to degree at most d.

Proof. An image of degree greater than d requires a surviving monomial using at least $d + 1$ blocks. Each such monomial survives with probability at most $q^{-(d+1)}$, by Lemma 7.7. Take the union bound over distinct monomials. Cancellation cannot increase the maximum degree of the surviving images. □

Theorem 7.9 (Lifting to bounded bottom degree)**.** *Let π be a semantic $\Sigma\Pi\Sigma\Pi$-PC refutation of $\mathrm{Lift}_q(F)$ over $\mathbb{F}$. For every integer $d \geq 1$ such that $M_{\mathrm{bot}}(\pi) < q^{d+1}$, there is a semantic $\Sigma\Pi\Sigma\Pi^{[d]}$-PC refutation π' of F with*

$$S(\pi') \leq S(\pi). \tag{38}$$

The projection also does not increase the number of proof vertices or product-term occurrences. In particular, the theorem applies with

$$d = \max\{1, \lfloor \log_q M_{\mathrm{bot}}(\pi) \rfloor\}.$$

Proof. Apply Lemma 7.8 to all bottom polynomials of π, and choose one projection under which they all have degree at most d. Project each represented line by keeping its outer sum and products:

$$P' = \sum_{a=1}^{t} c_a \prod_{b=1}^{r_a} \rho_A(f_{a,b}). \tag{39}$$

Write each projected bottom factor in its multilinear form; delete zero terms if desired.

Each target top product is the image under the same homomorphism of its source top product. If two source products are equal in the Boolean quotient, their images are equal as well. Zero source products remain zero. Therefore the number of distinct non-zero products cannot increase, proving (38).

Every semantic inference remains sound. If its projected premises vanish at y, its source premises vanish at $X_A(y)$ by (36). Source soundness forces the conclusion to vanish there, so its projected conclusion vanishes at y. The number of premises is unchanged.

By Lemma 7.6, initial equations project to input equations of F or to $0 = 0$, and the final equation remains $1 = 0$. Retain the projected representations of initial equations as well: there is no need to replace them by new canonical product representations. Thus no additional vertices or product terms are introduced.

For the last assertion, the displayed choice of d satisfies $M_{\mathrm{bot}}(\pi) < q^{d+1}$, including when $M_{\mathrm{bot}}(\pi)$ is an exact power of q. □

The theorem does not claim to preserve the number of bottom monomials: a surviving literal $1 - y_i$ may expand into two monomials. This does not affect the distinct-product lower bounds used here. If the source bottom degree is already at most b, affine substitution and multilinearisation give target bottom degree at most $\min\{b, d\}$.

Corollary 7.10 (Transfer of product-term lower bounds)**.** *Let F be any unsatisfiable CNF, let $q \geq 2$ and $d \geq 1$ be integers, and fix a field $\mathbb{F}$. Let $B_F(d)$ be a lower bound on the number of distinct non-zero top product terms in every semantic $\Sigma\Pi\Sigma\Pi^{[d]}$-PC refutation of F over $\mathbb{F}$. Every semantic $\Sigma\Pi\Sigma\Pi$-PC refutation π of $\mathrm{Lift}_q(F)$ satisfying*

$$M_{\mathrm{bot}}(\pi) < q^{d+1}$$

has

$$S(\pi) \geq B_F(d). \tag{40}$$

For every semantic $\Sigma\Pi\Sigma\Pi$-PC refutation π of $\mathrm{Lift}_q(F)$,

$$|\pi|_4 \geq \min\{q^{d+1}, B_F(d)\}. \tag{41}$$

If the base bound is available for $1 \leq d \leq d_0$, then

$$|\pi|_4 \geq \max_{1 \leq d \leq d_0} \min\{q^{d+1}, B_F(d)\}.$$

Proof. When $M_{\mathrm{bot}}(\pi) < q^{d+1}$, apply Theorem 7.9 to obtain $S(\pi) \geq S(\pi') \geq B_F(d)$. Otherwise (33) gives $|\pi|_4 \geq M_{\mathrm{bot}}(\pi) \geq q^{d+1}$. This proves the size bound; taking the maximum proves the last assertion. □

The circuit-size conclusion includes the bottom representation. Without a bound on $M_{\mathrm{bot}}(\pi)$, it does not give a lower bound on $S(\pi)$ alone for unrestricted bottom degree.

Corollary 7.11 (Quasipolynomial lower bounds from arbitrary base formulas)**.** *Let F be an unsatisfiable CNF on $n \geq 2$ variables, and let $d \geq 1$ be an integer. Suppose that every semantic $\Sigma\Pi\Sigma\Pi^{[d]}$-PC refutation of F over $\mathbb{F}$ has at least n^{d+1} distinct non-zero top product terms. Then every semantic $\Sigma\Pi\Sigma\Pi$-PC refutation π of $\mathrm{Lift}_n(F)$ over $\mathbb{F}$ satisfies*

$$|\pi|_4 \geq n^{d+1}.$$

In particular, if a family (F_n) of polynomial-size CNFs satisfies this hypothesis with $d = d(n) = \Omega(\log n)$, then $\mathrm{Lift}_n(F_n)$ has polynomial size and requires semantic $\Sigma\Pi\Sigma\Pi$-PC refutations of circuit size $n^{\Omega(\log n)}$.

Proof. Apply Corollary 7.10 with $q = n$ and $B_F(d) \geq n^{d+1}$. The formula-size claim follows from (34). □

7.3.4 Application to random $O(\log n)$-CNFs

We apply the general lifting consequence to the random CNF formulas studied in Section 7.1. Recall that F_n denotes a random ℓ-CNF formula on n variables with $m = n2^\ell$ clauses, where $\ell = \Theta(\log n)$ is chosen as in that subsection and every clause is sampled independently and uniformly from the clauses of width exactly ℓ. By Theorem 7.2, with high probability over the choice of F_n we have the required lower bounds for semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutations of F_n.

Relabel the variables of F_n as $y_1, \dots, y_n$, and put

$$G_n = \mathrm{Lift}_n(F_n).$$

The only randomness in G_n comes from the base formula F_n; the selector and exclusion clauses are deterministic.

Theorem 7.12 (Quasipolynomial lower bound for semantic $\Sigma\Pi\Sigma\Pi$-PC)**.** *The formula G_n has $2n^2$ variables, $O(n^3)$ clauses, and $O(n^3 \log n)$ literal occurrences. With probability $1 - o(1)$ it is unsatisfiable, and every semantic $\Sigma\Pi\Sigma\Pi$-PC refutation π of G_n over $\mathbb{F}_2$ has circuit size*

$$|\pi|_4 \geq n^{\Omega(\log n)} = \exp(\Omega((\log n)^2)).$$

This is a superpolynomial lower bound in its encoding length.

Proof. Fix the event in Theorem 7.2 and put $d = \lfloor c \log_2 n \rfloor$. For sufficiently large n, $d \geq 1$, and the base bound gives

$$B_{F_n}(d) \geq \frac{1}{4} \exp\left(\frac{an}{dC^d}\right) \geq \frac{1}{4} \exp\left(\frac{a\sqrt{n}}{d}\right).$$

Its logarithm grows faster than $(d+1)\ln n$, so $B_{F_n}(d) > n^{d+1}$ for sufficiently large n. Apply Corollary 7.11 to obtain

$$|\pi|_4 \geq n^{d+1} = n^{\Omega(\log n)}.$$

Since $q = n$, (34) gives $2n^2$ variables and $O(n^3)$ clauses. Moreover, F_n has $m = n2^\ell = \Theta(n^2)$ clauses and $\Lambda = m\ell = O(n^2 \log n)$ literal occurrences. Each literal occurrence in F_n produces n literal occurrences in the lifted clauses, while the selector and exclusion clauses contribute $O(n^3)$ further occurrences. Hence $\mathrm{Lift}_n(F_n)$ has $O(n^3 \log n)$ literal occurrences and therefore polynomial encoding length. □

7.4 Sparsification: PHP in semantic $\Sigma\Pi\Sigma\Pi$-PC

We now remove the bound on the bottom fan-in for PHP by sparsification. Starting with a purported small $\Sigma\Pi\Sigma\Pi$-PC refutation of PHP_n^{2n}, we sample a graph G using the construction in Lemma 7.4 and set every variable corresponding to a non-edge of G to zero. With high probability, G has the expansion and degree properties stated in Lemma 7.4. The same restriction also eliminates every bottom product gate of fan-in greater than w, provided that the original circuit is sufficiently small. The refutation then becomes a $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutation of $\mathsf{PHP}(G)$, and Theorem 7.5 applies.

Lemma 7.13 (Sparsification to an expander)**.** *Let π be a semantic $\Sigma\Pi\Sigma\Pi$-PC refutation of PHP_n^{2n} of total circuit size M. Sample G as in Lemma 7.4, with Δ as defined there. For every integer $w \geq 1$, with probability at least*

$$1 - M\left(\frac{\Delta}{n}\right)^{w+1} - o(1),$$

the graph G satisfies the conclusions of Lemma 7.4. On the same event, restricting π to G means applying to every proof line the partial assignment

$$\rho_G(x_{u,v}) = \begin{cases} x_{u,v}, & (u,v) \in E(G), \\ 0, & (u,v) \notin E(G). \end{cases}$$

The resulting proof is a semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutation of $\mathsf{PHP}(G)$ on the hole domain

$$\Omega_G = \left\{\alpha \in \{0,1\}^{E(G)} : \alpha_{u,v}\alpha_{u',v} = 0 \text{ for all distinct } u, u' \text{ with } (u,v),(u',v) \in E(G)\right\}.$$

That is, every semantic inference is required to be sound only on assignments in Ω_G. Every initial line of the restricted proof, viewed as a function on Ω_G, is equal either to an axiom of $\mathsf{PHP}(G)$ or to the zero polynomial. The restriction does not increase circuit size, so the resulting refutation has circuit size at most M.

Proof. First multilinearise each bottom product term and discard those that contain two distinct variables associated with the same hole. These operations preserve every line on the hole domain and do not increase circuit size. Every remaining bottom product term mentions each hole at most once. We sample G as in Lemma 7.4 and set $x_{u,v} = 0$ whenever (u, v) is not an edge of G. Every bottom product term t of fan-in greater than w mentions at least $w+1$ holes. Thus, t survives with probability at most

$$\left(\frac{\Delta}{n}\right)^{w+1}.$$

Since there are at most M such bottom product gates, a union bound shows that all of them are eliminated with probability at least

$$1 - M\left(\frac{\Delta}{n}\right)^{w+1}.$$

By Lemma 7.4, the probability that G fails to have the expansion and degree properties stated there is $o(1)$. Thus, both events occur with probability at least

$$1 - M\left(\frac{\Delta}{n}\right)^{w+1} - o(1)$$

as claimed. Semantic refutations are closed under partial restrictions, so this restriction turns PHP_n^{2n} into $\mathsf{PHP}(G)$ and simultaneously transforms π into a semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutation of $\mathsf{PHP}(G)$ on its hole domain, without increasing its circuit size. □

Theorem 7.14 (Complete unary PHP)**.** *For every fixed $0 < \gamma < 1/2$ and all sufficiently large n, every semantic $\Sigma\Pi\Sigma\Pi$-PC refutation of PHP_n^{2n} has total circuit size at least*

$$\exp\left(\gamma\frac{(\ln n)^2}{\ln\ln n}\right).$$

Proof. Let M be the total circuit size of a refutation of PHP_n^{2n}.

Fix a constant β with $\gamma < \beta < 1/2$, and suppose for a contradiction that

$$M < \exp\left(\gamma\frac{(\ln n)^2}{\ln\ln n}\right).$$

Put

$$w = \left\lfloor \beta\frac{\ln n}{\ln\ln n}\right\rfloor.$$

Put $\Delta = \lceil 24\ln(2n)\rceil$. We have

$$\ln\left(M\left(\frac{\Delta}{n}\right)^{w+1}\right) \le (\gamma - \beta + o(1))\frac{(\ln n)^2}{\ln\ln n} < 0.$$

Thus, by our choice of M, the total failure probability in Lemma 7.13 is less than 1 for all sufficiently large n. We may therefore fix a graph G satisfying Lemma 7.4 such that the restriction of the original refutation is a semantic $\Sigma\Pi\Sigma\Pi^{[w]}$-PC refutation of $\mathsf{PHP}(G)$ on its hole domain.

Pad each hole class as in (29). Since the restricted proof does not depend on these variables, it remains sound on the domain of $\mathcal{F}_0$. We may therefore apply Theorem 7.5. The number of distinct non-zero top product terms is at most the circuit size M, while the theorem gives

$$M \ge \frac{1}{4}\exp\left(\frac{c_0 n}{w\Delta(C_0\Delta^2)^w}\right) = \exp\left(n^{1-2\beta-o(1)}\right).$$

Here we used

$$\ln(w\Delta) + w\ln(C_0\Delta^2) = 2\beta\ln n + o(\ln n).$$

This contradicts the assumed upper bound on M and proves the theorem. □

The formula PHP_n^{2n} has size $\Theta(n^3)$, so Theorem 7.14 is superpolynomial in the input size.

Consequences for $\mathsf{Res}(\log)$. We recall the relevant extension of Resolution by bounded-width conjunctions. A k-term is a conjunction of at most k literals, and a k-DNF is an unbounded disjunction of k-terms. The proof system $\mathsf{Res}(k)$ operates with k-DNFs and uses the standard weakening, conjunction-introduction, and cut rules; see, for example, [7]. We measure the size of a refutation by the total number of symbols in all its lines. A refutation of size S using these rules is called a $\mathsf{Res}(\log)$ refutation if every term occurring in it has width at most $\lceil \log_2 S \rceil$.

Proposition 7.15 (Simulation of $\mathsf{Res}(\log)$)**.** *Let φ be a CNF formula. If φ has a $\mathsf{Res}(\log)$ refutation of size S, then its polynomial encoding $\mathcal{F}_\varphi$ has a semantic $\Sigma\Pi\Sigma\Pi$-PC refutation over $\mathbb{F}_2$ of total circuit size $O(S^2)$.*

Proof. Consider a proof line

$$D = T_1 \vee \cdots \vee T_m,$$

where every T_i is a term of width at most $\lceil \log_2 S \rceil$. Associate with T_i its indicator polynomial

$$\tau_i = \prod_{x_j \in T_i} x_j \prod_{\neg x_j \in T_i} (1 - x_j).$$

On the Boolean cube, $\tau_i = 1$ precisely when T_i is satisfied. We encode D by

$$[\![D]\!] = \prod_{i=1}^{m} (1 - \tau_i).$$

Thus, for every Boolean assignment a,

$$[\![D]\!](a) = 0 \quad \Longleftrightarrow \quad a \models D.$$

Expanding each τ_i separately yields a sparse polynomial with at most

$$2^{|T_i|} \le 2^{\lceil \log_2 S \rceil} = O(S)$$

monomials. We keep the outer product in $[\![D]\!]$ unexpanded. Consequently, $[\![D]\!]$ has a $\Pi\Sigma\Pi$ representation, and hence a $\Sigma\Pi\Sigma\Pi$ representation with a single top product term, of size $O(S|D|)$.

Every inference in $\mathsf{Res}(\log)$ has at most two premises and is semantically sound. The equivalence above therefore translates it into a valid semantic $\Sigma\Pi\Sigma\Pi$-PC inference. Initial clauses translate into their standard polynomial encodings, while the empty DNF translates into the empty product 1. We thus obtain a semantic $\Sigma\Pi\Sigma\Pi$-PC refutation. Since

$$\sum_{D \in \pi} |D| \le S,$$

its total circuit size is

$$O\left(S \sum_{D \in \pi} |D| \right) = O(S^2).$$

□

Corollary 7.16 ($\mathsf{Res}(\log)$ lower bound for weak PHP)**.** *For every fixed $0 < \delta < 1/4$ and all sufficiently large n, every $\mathsf{Res}(\log)$ refutation of PHP_n^{2n} has size*

$$S \ge \exp\left(\delta \frac{(\ln n)^2}{\ln \ln n} \right).$$

Proof. Choose a constant γ such that

$$2\delta < \gamma < \frac{1}{2}.$$

By Proposition 7.15, a Res(log) refutation of size S yields a semantic $\Sigma\Pi\Sigma\Pi$-PC refutation of total circuit size at most CS^2 for some constant C. On the other hand, Theorem 7.14 gives

$$CS^2 \geq \exp\left(\gamma \frac{(\ln n)^2}{\ln \ln n}\right).$$

It follows that

$$\ln S \geq \frac{\gamma}{2} \frac{(\ln n)^2}{\ln \ln n} - O(1) \geq \delta \frac{(\ln n)^2}{\ln \ln n}$$

for all sufficiently large n. □

The result of Maciel, Pitassi, and Woods [53] implies that the weak pigeonhole principle has quasipolynomial-size Res(log) refutations; see also [7]. Together with Corollary 7.16, this shows that both the known upper bound and our lower bound for Res(log) are quasipolynomial, although their exponents do not match. In particular, the quasipolynomial upper bound cannot be improved to a polynomial one in Res(log). The broader question of whether the weak pigeonhole principle has polynomial-size constant-depth Frege proofs remains open.

8 Lower Bounds for Semantic $\mathsf{AC}^0[2]$-Frege

Fix $d \geq 2$ and an unsatisfiable CNF F_n of constant initial width w_0 with minimum PC refutation degree $\Delta_n \geq n^C$ for some positive constant C. We turn this degree lower bound into a size lower bound for semantic depth-d $\mathsf{AC}^0[2]$-Frege by replacing the literals of F_n with small circuit gadgets. Our goal is to apply Corollary 4.18: we need a large space U of multipliers admitting non-zero relative pseudo-reduction operators, and parameters for which the exponent θ is large. More precisely, we keep the dimension ratio $\dim(U)/H_{nm}(k)$ close to one, where nm is the number of variables after lifting.

Let g be a Boolean circuit in m variables and define $g_0(X) = g(X)$ and $g_1(X) = g(\overline{X})$, where $\overline{X}$ is the coordinatewise complement of X. For each variable z_i of F_n, introduce a block X_i of fresh variables and replace the literals z_i and $\neg z_i$ by $g_1(X_i)$ and $g_0(X_i)$, respectively. We choose an AND-rooted gadget of depth at most $d+1$ and use distributivity to obtain lifted axioms of depth at most d. The details are given in Section 8.1.

We require the disjointness property $g_0(X) \wedge g_1(X) = 0$: the gadgets replacing z_i and $\neg z_i$ cannot both evaluate to 1 on the same assignment to X_i. This ensures that any satisfying assignment to the lifted axioms would give a satisfying assignment to F_n, and hence the lifting preserves unsatisfiability; see Proposition 8.3.

We also require g to be antitone. This lets us construct a multiplier space U of sufficiently large dimension such that a derivation of any non-zero $f \in U$ from the lifted axioms multiplied by f, using lines fQ with Q of low degree, yields a low-degree PC refutation of F_n. The degree lower bound for F_n rules out such derivations, giving the non-zero relative pseudo-reduction operators required by the corollary. We prove this in Section 8.2.

The space U is spanned by low-degree monomials whose support in every block is accepted by g. Antitonicity makes the family of accepted supports downward closed. The corresponding monomials then form a common basis for functions on the two sets where g_0 and g_1 are true. Expanding in this basis lets us transfer derivations to the base variables: the degree of each coefficient of fQ is controlled by the degree of Q, independently of the degree of f.

To make U large, we want almost all monomials up to the chosen degree cutoff to have accepted supports in every block. The complement of majority on an odd number of variables accepts every support of size less than $m/2$. For a cutoff slightly below $nm/2$ and sufficiently large blocks, almost all monomial supports satisfy this condition in every block. This is the property that our approximate majority gadget must retain, allowing a small fraction of rejected supports.

The complement of majority satisfies the required structural and dimension properties. However, computing it exactly at the required depth would make the gadget too large for our lower bound to

be superpolynomial in the size of the lifted instance. We therefore construct a smaller approximate majority gadget using the ϵ-coin problem, for which Limaye et al. [52] proved tight circuit size bounds; see Section 8.3. For our parameters, the gadget has size polynomial in n and gives $\dim(U)/H_{nm}(k) = 1 - o(1)$. Finally, Section 8.4 chooses the approximation parameters to make θ large while keeping the degree of the proof-line approximations below Δ_n. The resulting bound in Theorem 8.1 is superpolynomial in the size of the lifted instance.

We write $\mathrm{Lift}_g(F_n)$ for the result of lifting F_n with the gadget g.

Theorem 8.1 (Size lower bound for $\mathrm{Lift}_g(F_n)$)**.** *Fix $d \geq 2$, a constant w_0, and $C > 0$. Let F_n be an unsatisfiable CNF on n variables of maximum clause width at most w_0, with $\Delta_n = \deg_{\mathsf{PC}}(F_n) \geq n^C$. There exists a gadget g such that $\mathrm{Lift}_g(F_n)$ has total circuit size $N = n^{O_{d,w_0}(1)}$, and every semantic depth-d $\mathsf{AC}^0[2]$-Frege refutation of $\mathrm{Lift}_g(F_n)$ has size*

$$\exp\Big(\Omega_d\Big(\ln n\,(\ln \Delta_n)^{1/(d-1)}\Big)\Big).$$

The formula has a polynomial-size semantic depth-$(d+1)$ AC^0-Frege refutation.

We construct the gadget g in Section 8.3 and choose its parameters in (52).

Applying the theorem to random 3-CNFs at a sufficiently large fixed clause density gives the families asserted in the introduction. By Lemma 6.1, these formulas are unsatisfiable with probability $1 - o(1)$, and Lemmas 6.2 and 5.3 give $\Delta_n = \Omega(n)$ with probability $1 - o(1)$. The lifted family has total circuit size $N = n^{O_d(1)}$, so the lower bound becomes $\exp\Big(\Omega_d\Big((\log N)^{d/(d-1)}\Big)\Big)$.

8.1 The lifting operation

Fix $d \geq 2$ and an AND/OR circuit g on m variables with the following properties.

1. The depth of g is at most $d + 1$.
2. g is AND-rooted.
3. g is antitone.
4. For every $X \in \{0,1\}^m$, $g(X) \wedge g(\overline{X}) = 0$.

Let $g_0(X) = g(X)$ and $g_1(X) = g(\overline{X})$. Observe that for every $X \in \{0,1\}^m$,

$$g_0(X) \wedge g_1(X) = 0. \tag{42}$$

Let F_n be a CNF in variables $z_1, \ldots, z_n$, with initial width at most $w_0 = O(1)$. Write $z_i^1 = z_i$ and $z_i^0 = \neg z_i$, so z_i^b is the literal asserting $z_i = b$. For every $i \in [n]$, we introduce a block of m fresh variables $X_i = (x_{i,1}, \ldots, x_{i,m})$ and replace each literal z_i^b with $g_b(X_i)$. This replaces each clause

$$C = \bigvee_{(i,b)\in I} z_i^b, \qquad I \subseteq [n] \times \{0,1\}, \quad 1 \leq |I| \leq w_0$$

with

$$\bigvee_{(i,b)\in I} g_b(X_i),$$

which is a circuit of depth at most $d + 2$.

Observe that these replacements might not be complementary: both $g_0(X_i)$ and $g_1(X_i)$ can be zero, but (42) ensures that they cannot both be one.

Since g is AND-rooted, we can write for each $b \in \{0,1\}$,

$$g_b(X) = \bigwedge_{j\in J} H_{b,j}(X),$$

where J is nonempty and each $H_{b,j}$ is an OR-rooted circuit of depth at most d (we can add a rudimentary arity-1 OR gate if $H_{b,j}$ is a literal or a constant).

For a clause $C = \bigvee_{(i,b)\in I} z_i^b$, let $\alpha \in J^I$ choose one index for each literal. Distributivity gives

$$\bigvee_{(i,b)\in I} g_b(X_i) = \bigvee_{(i,b)\in I} \bigwedge_{j\in J} H_{b,j}(X_i) = \bigwedge_{\alpha\in J^I} \bigvee_{(i,b)\in I} H_{b,\alpha(i,b)}(X_i). \tag{43}$$

Definition 8.2 (Lifting operation)**.** The lifted formula $\mathrm{Lift}_g(F_n)$ is defined as

$$\mathrm{Lift}_g(F_n) := \{A_{C,\alpha} : C = \bigvee_{(i,b)\in I} z_i^b \in F_n,\ \alpha \in J^I\}, \tag{44}$$

where

$$A_{C,\alpha}(X) := \bigvee_{(i,b)\in I} H_{b,\alpha(i,b)}(X_i).$$

Each circuit $A_{C,\alpha}(X)$ has depth at most d, since its OR root merges with the OR roots of its children. It depends on at most w_0 blocks.

For a size bound, let S_g be the circuit size of the gadget g. Each clause contributes at most $|J|^{w_0}$ axioms, each of size $O(w_0 S_g)$. Hence the total size of $\mathrm{Lift}_g(F_n)$ is at most $|F_n||J|^{w_0}O(w_0 S_g)$. In particular, for constant w_0 and a polynomial-size gadget, the lifted family has size polynomial in n.

We need to show that this replacement preserves unsatisfiability. This is guaranteed by (42).

Proposition 8.3. *If F_n is unsatisfiable, then $\mathrm{Lift}_g(F_n)$ is unsatisfiable.*

Proof. If all $A_{C,\alpha}$ were true, (43) would make every substituted base clause true. Set $z_i = b$ whenever $g_b(X_i) = 1$, and assign z_i arbitrarily when $g_0(X_i) = g_1(X_i) = 0$. These choices are consistent by (42). Every true substituted literal then gives a true base literal, so every clause of F_n would be satisfied, which is a contradiction. □

8.1.1 Why the gadget needs the extra depth

Enumerate the circuits in $\mathrm{Lift}_g(F_n)$ as $A_1, \ldots, A_h$. By construction, their depth is at most d. Now set $D_i = \bigwedge_{j\le i} A_j$ for each $i \in [h]$. Each D_i has depth at most $d+1$, and D_{i+1} follows from D_i and A_{i+1} by a binary semantic inference. Unsatisfiability makes D_h identically zero, so one further inference gives the constant-zero circuit. This refutation has size $O(S + h)$, where S is the circuit size of $\mathrm{Lift}_g(F_n)$: each conjunction D_i introduces at most one new circuit. Moreover, this refutation only uses AND and OR gates.

If the gadget g had depth at most d, each circuit in $\mathrm{Lift}_g(F_n)$ would have depth at most $d-1$, and this same conjunction argument would yield a polynomial-size depth-d refutation of $\mathrm{Lift}_g(F_n)$. The extra gadget layer is therefore needed to prevent this refutation within the lower-bound depth, while allowing it at depth $d+1$.

8.2 Degree bounds on the lifted formula

Let g satisfy the assumptions of Section 8.1, with $g_0 = g$ antitone and $g_1(X) = g_0(\overline{X})$, where $\overline{X}$ denotes coordinatewise complementation. Put $v = nm$ and $\mathcal{S} = \{[\![A]\!] : A \in \mathrm{Lift}_g(F_n)\}$. We construct a space U and prove nonvanishing of its naïve relative pseudo-reductions by transferring coefficients to the base variables.

The multiplier subspace. Let $\mathcal{L} = \{S \subseteq [m] : g_0(1_S) = 1\}$, where 1_S is the indicator vector of S. This family is downward closed because g_0 is antitone. For $S = (S_1, \ldots, S_n) \in \mathcal{L}^n$, put

$$x_S = \prod_{i=1}^{n} \prod_{j\in S_i} x_{i,j},$$

and define the space

$$U = \mathrm{span}\left\{ x_S : S \in \mathcal{L}^n,\ \sum_i |S_i| \le k \right\}. \tag{45}$$

The main result of this subsection is the following lemma, which shows that $\mathcal{R}_f(f) \neq 0$ for every non-zero $f \in U$.

Lemma 8.4 (Degree transfer and relative pseudo-reductions)**.** *Let F_n be an unsatisfiable CNF, let $0 \neq f \in U$, and let $\ell \geq 0$ be an integer. If f has a* PC *derivation from $f\mathcal{S}$ entirely in $f\mathcal{B}_v^{\leq \ell}$, then F_n has a* PC *refutation of degree at most $\ell + 1$. Consequently, if every base refutation has degree at least Δ_n and $\ell + 1 < \Delta_n$, the naïve relative pseudo-reduction $\mathcal{R}_{f,\ell}$ for $\mathcal{S}$ is non-zero.*

Taking $\ell = 3q$ supplies the operators required by Theorem 4.16 whenever $3q + 1 < \Delta_n$. We prove the lemma after constructing the coefficient maps.

A common basis and its coefficients. The gadgets g_0 and g_1 may both vanish simultaneously, so we restrict to inputs on which exactly one is true. Define

$$\Omega_0 = \{X : g_0(X) = 1\} = \{1_S : S \in \mathcal{L}\},$$
$$\Omega_1 = \{X : g_1(X) = 1\} = \{1_{[m]\setminus S} : S \in \mathcal{L}\}$$

and observe that disjointness (42) implies that $\Omega_0 \cap \Omega_1 = \varnothing$.

For $z \in \{0,1\}^n$, put $\Omega(z) = \prod_i \Omega_{z_i}$. These fibres partition $\Omega = (\Omega_0 \cup \Omega_1)^n$, and every substituted literal $g_b(X_i)$ equals the base literal z_i^b on $\Omega(z)$.

Lemma 8.5 (Common monomial basis)**.** *The monomials $x_S = \prod_{j\in S} x_j$, $S \in \mathcal{L}$ give a basis of both $\mathbb{F}_2^{\Omega_0}$ and $\mathbb{F}_2^{\Omega_1}$.*

Proof. Define the evaluation matrix on Ω_0 by

$$\mathsf{eval}_{\Omega_0} := \big(x_S(1_T)\big)_{S,T\in\mathcal{L}}.$$

Its rows are indexed by the monomials x_S, and its columns by the points $1_T \in \Omega_0$. Since x_S is the product of the variables indexed by S, its value at 1_T is 1 exactly when every element of S belongs to T. Thus

$$x_S(1_T) = \begin{cases} 1 & \text{if } S \subseteq T, \\ 0 & \text{otherwise.} \end{cases}$$

Order the rows and columns by non-decreasing cardinality, using the same order for sets of equal size. A non-zero entry requires $|S| \leq |T|$, and if $|S| = |T|$, it requires $S = T$. Hence eval_{Ω_0} is upper triangular with every diagonal entry equal to 1, so it is invertible. Its rows are therefore linearly independent. Since $|\Omega_0| = |\mathcal{L}|$, the monomials x_S, $S \in \mathcal{L}$, form a basis of $\mathbb{F}_2^{\Omega_0}$.

The coordinatewise complementation $\sigma\colon x_j \mapsto 1 + x_j$ is a bijection from Ω_0 to Ω_1. Under this bijection, the function x_S on Ω_1 corresponds to the function

$$\sigma(x_S) = \prod_{j\in S}(1 + x_j) = \sum_{T\subseteq S} x_T$$

on Ω_0. Every $T \subseteq S$ belongs to $\mathcal{L}$ because $\mathcal{L}$ is downward closed. The sum contains x_S with coefficient 1, and all its other terms have smaller degree. Consequently, these functions are obtained from the basis $\{x_S : S \in \mathcal{L}\}$ by a triangular change of basis with diagonal entries equal to 1. They therefore also form a basis of $\mathbb{F}_2^{\Omega_0}$. Transporting this basis along the complementation bijection shows that the monomials x_S, $S \in \mathcal{L}$, form a basis of $\mathbb{F}_2^{\Omega_1}$. □

Let $\zeta\colon \Omega \to \{0,1\}^n$ send each point to its fibre label: for every $z \in \{0,1\}^n$ and $X \in \Omega(z)$, $\zeta(X) = z$. A base polynomial $b \in \mathcal{B}_n$ acts on functions on Ω by $(b \cdot a)(X) = b(\zeta(X))a(X)$: it multiplies each fibre $\Omega(z)$ by the scalar $b(z)$. On each fibre $\Omega(z)$, the monomials x_S, $S \in \mathcal{L}^n$, form a basis of $\mathbb{F}_2^{\Omega(z)}$. The unique coefficients in these expansions are functions of z, hence elements of $\mathcal{B}_n$. This gives the decomposition

$$\mathbb{F}_2^{\Omega} = \bigoplus_{S\in\mathcal{L}^n} \mathcal{B}_n\, x_S. \tag{46}$$

For $u \in \mathcal{B}_v$, let $c_S(u) \in \mathcal{B}_n$ be the coordinates of its restriction to Ω. Thus

$$u(X) = \sum_{S\in\mathcal{L}^n} c_S(u)(z)x_S(X) \qquad \text{for every } X \in \Omega(z).$$

The coordinate maps are $\mathcal{B}_n$-linear on functions on Ω. In particular, if $f = \sum_S \beta_S x_S \in U$, then $c_S(f) = \beta_S$, so a nonzero f has a coordinate equal to 1.

A low-degree polynomial can involve many blocks, but each of its monomials mentions few. This locality controls the degrees of its coefficients.

Lemma 8.6 (Coefficient degree and locality)**.** *Let $f \in \operatorname{span}\{x_T : T \in \mathcal{L}^n\}$ and $p \in \mathcal{B}_v$. Then*

$$\deg c_S(fp) \le \deg p \qquad (S \in \mathcal{L}^n).$$

If p involves only blocks indexed by $J \subseteq [n]$, then $c_S(fp)$ depends only on z_J.

Proof. By linearity in f and p, it suffices to take $f = x_T$, where $T \in \mathcal{L}^n$, and p a monomial. Let J be the blocks mentioned by p. Expand px_T on $\Omega(z)$ block by block. For $i \notin J$, the polynomial p contains no variables from block X_i, so the factor in this block is x_{T_i}. By Lemma 8.5, this monomial is a basis element on both Ω_0 and Ω_1. Its coefficient at x_{S_i} is therefore 1 if $S_i = T_i$ and 0 otherwise, independently of z_i. Thus $c_S(px_T)(z)$ depends only on z_J. Writing its value at $z_J = a$ as η_a, interpolation gives

$$c_S(px_T)(z) = \sum_{a \in \{0,1\}^J} \eta_a \chi_{J,a}(z), \qquad \chi_{J,a}(z) = \prod_{i \in J : a_i = 1} z_i \prod_{i \in J : a_i = 0} (1 + z_i).$$

Here $\chi_{J,a}$ indicates the assignment $z_J = a$, so every summand has degree at most $|J| \le \deg p$. Expanding arbitrary f and p into monomials proves both assertions. □

We can now prove Lemma 8.4.

Proof of Lemma 8.4. Choose representations $fQ_1, \dots, fQ_L$ of the derivation lines with $\deg Q_i \le \ell$, taking $Q_L = 1$. We derive all their coordinates $c_S(fQ_i)$ in the base variables. By Lemma 8.6, each coordinate has degree at most ℓ.

For an input $fh_{C,\alpha}$, where $h_{C,\alpha} = [\![A_{C,\alpha}]\!]$, each coordinate depends only on the variables of C. If z satisfies C, all its substituted literals have their base values on $\Omega(z)$, so the substituted clause and all its distributed conjuncts are true. Hence $h_{C,\alpha}$ vanishes on that fibre, as do all coordinates of $fh_{C,\alpha}$. Each coordinate is therefore supported on the unique falsifying assignment of C, and is either zero or $[\![C]\!]$. Thus initial coordinates are base axioms or zero, and every base axiom used has degree at most ℓ.

Suppose a proof line follows by multiplying fp by $x_{i,j}$. On Ω, expanding in the common basis gives

$$\begin{aligned} fx_{i,j}p &= x_{i,j} \sum_{T \in \mathcal{L}^n} c_T(fp) x_T \\ &= \sum_{T \in \mathcal{L}^n} c_T(fp) \sum_{S \in \mathcal{L}^n} c_S(x_{i,j} x_T) x_S \\ &= \sum_{S \in \mathcal{L}^n} \left(\sum_{T \in \mathcal{L}^n} c_T(fp) c_S(x_{i,j} x_T) \right) x_S. \end{aligned}$$

Comparing coefficients yields

$$c_S(fx_{i,j}p) = \sum_{T \in \mathcal{L}^n} c_T(fp) c_S(x_{i,j} x_T). \tag{47}$$

Multiplication by $x_{i,j}$ changes only block i; in every other block, the factor remains a common basis monomial on either fibre. Thus, by Lemma 8.6 with $f = x_T$ and $p = x_{i,j}$, each multiplying coefficient $c_S(x_{i,j}x_T)$ depends only on z_i, and is an affine form in z_i. Each summand can therefore be derived from $c_T(fp)$ by variable multiplication and addition in degree at most $\ell + 1$. Adding the summands preserves this bound. The resulting coordinates themselves have degree at most ℓ, by the new line's representation in $f\mathcal{B}_v^{\le \ell}$.

At the final line f, some coordinate is 1 since $f \neq 0$. Its derivation is a base refutation of degree at most $\ell + 1$. If $\ell + 1 < \Delta_n$, no such derivation of f exists, so the quotient map (9) is non-zero. □

8.3 Gadget construction

We construct the gadget satisfying the structural requirements of Section 8.1.

Lemma 8.7 (Gadgets with a dimension bound)**.** *Fix $d \geq 2$ and a sufficiently small $0 < \epsilon \leq 1/256$. Put $L := \epsilon^{-1/(d-1)}$ and let m be the smallest odd integer at least $128L^{2d-1}$. There exists an AND-rooted gadget $g\colon \{0,1\}^m \to \{0,1\}$ represented by an AND/OR circuit of depth at most $d+1$ and size $\exp(O_d(L))$, such that g is antitone and $g(X) \wedge g(\overline{X}) = 0$ for every $X \in \{0,1\}^m$.*

For this gadget, define U as in (45). For every $n \geq 1$, put $v = nm$. Then, for every integer $0 \leq k \leq (1/2 - \epsilon)v$,

$$\gamma = \frac{\dim(U)}{H_v(k)} \geq 1 - (v+1)ne^{-\epsilon^{-1/(d-1)}}. \tag{48}$$

A covering gadget with biased rarity. We obtain the gadget g from a monotone formula C_m by $g(X) = \neg C_m(X)$. Write $\mathrm{Bern}(p)^m$ for the distribution of m independent bits, each equal to 1 with probability p.

Definition 8.8 (Coin problem and distinguishers)**.** For $0 < \epsilon < 1$, the ϵ-coin problem is to distinguish independent samples of a coin of bias $(1-\epsilon)/2$ from independent samples of a coin of bias $(1+\epsilon)/2$. A Boolean function $T\colon \{0,1\}^m \to \{0,1\}$ is a distinguisher for the ϵ-coin problem with error at most η if

$$\Pr_{x \sim \mathrm{Bern}((1-\epsilon)/2)^m}[T(x) = 1] \leq \eta \quad \text{and} \quad \Pr_{x \sim \mathrm{Bern}((1+\epsilon)/2)^m}[T(x) = 0] \leq \eta.$$

We use the following monotone formula upper bound.

Theorem 8.9 (Limaye et al. [52])**.** *Fix $d \geq 2$. For all sufficiently small $\epsilon > 0$, there is a monotone AND/OR formula of depth at most d and size $\exp(O_d(\epsilon^{-1/(d-1)}))$ solving the ϵ-coin problem with error at most $1/10$.*

Lemma 8.10 (A shifted coin distinguisher)**.** *Fix $d \geq 2$. For all sufficiently small $\epsilon > 0$, there exist an integer $m \geq 1$ and a monotone, AND-rooted formula T on m input variables, of depth at most d and size $\exp(O_d(\epsilon^{-1/(d-1)}))$, such that*

$$\Pr_{x \sim \mathrm{Bern}(1/2-\epsilon)^m}[T(x) = 1] \leq \tfrac{1}{5}, \qquad \Pr_{x \sim \mathrm{Bern}(1/2)^m}[T(x) = 0] \leq \tfrac{1}{5}.$$

Proof. Set $\epsilon' = \epsilon/(1+\epsilon)$, and let T_0 be the formula supplied by Theorem 8.9 with parameter ϵ'. Taking its monotone dual $1 - T_0(1-\cdot)$ if necessary, we may assume that T_0 is AND-rooted; duality interchanges the two error probabilities.

Independently for each distinct input variable of T_0, replace it by the constant 1 with probability ϵ', and retain it otherwise. Use the same replacement at every occurrence of a variable. Under product bias u, averaging over these replacements makes the original formula see independent inputs of bias

$$q(u) = \epsilon' + (1-\epsilon')u, \qquad q(1/2) = \frac{1+\epsilon'}{2}, \qquad q(1/2-\epsilon) = \frac{1-\epsilon'}{2}.$$

Thus the expected sum of the rejection probability at bias $1/2$ and the acceptance probability at bias $1/2 - \epsilon$ is at most $1/5$. Fix substitutions for which this sum is at most $1/5$. Each error is then at most $1/5$, and the resulting formula T is monotone and has the required depth and size, since $\epsilon' = \Theta(\epsilon)$. Its AND root may be retained when constant substitutions simplify some of its children. □

Lemma 8.11 (Amplification by conjunction)**.** *Fix $d \geq 2$. For all sufficiently small $\epsilon > 0$ and every integer $t \geq 1$, there exist an integer $m \geq 1$ and a monotone, AND-rooted formula T on m input variables, of depth at most d and size $t\exp(O_d(\epsilon^{-1/(d-1)}))$, such that*

$$\Pr_{x \sim \mathrm{Bern}(1/2)^m}[T(x) = 1] \geq (4/5)^t, \qquad \Pr_{x \sim \mathrm{Bern}(1/2-\epsilon)^m}[T(x) = 1] \leq (1/5)^t.$$

Proof. Let T' on m' input variables be supplied by Lemma 8.10. Take T to be the conjunction of t copies of T' on disjoint sets of variables, and put $m = tm'$. The acceptance events of the copies are independent under each product measure, so their probabilities multiply. The new conjunction merges with the AND roots of the copies, preserving the depth bound, and the size is $O(t|T'|)$. □

Lemma 8.12 (A one-sided covering gadget)**.** *Fix $d \geq 2$. For all sufficiently small $\epsilon > 0$, put $L = \epsilon^{-1/(d-1)}$ and let m be the smallest odd integer at least $128L^{2d-1}$. There exists a monotone, alternating, OR-rooted AND/OR formula C_m of depth at most $d+1$ and size $\exp(O_d(L))$ such that*

$$|x| \geq m/2 \implies C_m(x) = 1, \qquad \Pr_{X \sim \mathrm{Bern}(1/2-\epsilon)^m}[C_m(X) = 1] \leq e^{-L}.$$

Every child of the OR root is AND-rooted and has depth at most d. In particular, $m = \Theta_d(L^{2d-1})$.

Proof. Put $L = \epsilon^{-1/(d-1)}$ and $\epsilon' = 3\epsilon/4$. Apply Lemma 8.11 with parameter ϵ' and

$$t = \lceil 10L/\ln 4 \rceil.$$

The resulting AND-rooted formula, denoted by T, has depth at most d, size $t\exp(O_d(L))$, and, writing m' for its number of distinct input variables, satisfies

$$\Pr_{y \sim \mathrm{Bern}(1/2)^{m'}}[T(y) = 1] \geq (4/5)^t, \qquad \Pr_{y \sim \mathrm{Bern}(1/2-3\epsilon/4)^{m'}}[T(y) = 1] \leq (1/5)^t.$$

Sampling with replacement (Amano [6]). Introduce variables $x_1, \ldots, x_m$. Replace each variable of T by x_j, where j is chosen independently and uniformly at random from $[m]$. Use the *same* label at repeated occurrences of one formal variable. Let $T^\sigma(X)$ be the resulting formula. If a fixed $x \in \{0,1\}^m$ has density $u = |x|/m$, the values of the distinct formal variables under a random σ have joint distribution $\mathrm{Bern}(u)^{m'}$. Therefore

$$\Pr_\sigma[T^\sigma(x) = 1] = \Pr_{y \sim \mathrm{Bern}(u)^{m'}}[T(y) = 1].$$

Take $s = \lceil 4m(5/4)^t \rceil$ independent random relabellings $\sigma_1, \ldots, \sigma_s$, write $\sigma = (\sigma_1, \ldots, \sigma_s)$, and set

$$C_m(X) = \bigvee_{j=1}^{s} T^{\sigma_j}(X).$$

Coverage. For every fixed x of weight at least $m/2$, monotonicity and the acceptance bound for T at bias $1/2$ imply that a random copy accepts with probability at least $(4/5)^t$. Hence

$$\Pr_\sigma[C_m(x) = 0] \leq (1 - (4/5)^t)^s \leq e^{-4m}.$$

Taking a union bound over all such inputs,

$$\Pr_\sigma[\exists x,\ |x| \geq m/2 :\ C_m(x) = 0] \leq 2^m e^{-4m} < \tfrac{1}{2}. \tag{49}$$

Acceptance under the lower bias. Write $A(\sigma) = \Pr_{X \sim \mathrm{Bern}(1/2-\epsilon)^m}[C_m(X) = 1]$. For an x of weight at most $(1/2-3\epsilon/4)m$, a union bound and monotonicity give $\Pr_\sigma[C_m(x) = 1] \leq s5^{-t}$. Hoeffding's inequality yields

$$\Pr[|X| > (1/2 - 3\epsilon/4)m] \leq \exp(-m\epsilon^2/8) \leq e^{-16L}.$$

Splitting according to this weight threshold gives

$$\mathbb{E}_\sigma\, A(\sigma) \leq s5^{-t} + e^{-16L} = o(e^{-L}),$$

since $s5^{-t} \leq 4m4^{-t} + 5^{-t}$, $4^{-t} \leq e^{-10L}$, and $\ln m = O_d(\ln L) = o(L)$. By Markov's inequality,

$$\Pr_\sigma[A(\sigma) > e^{-L}] \leq e^L\, \mathbb{E}_\sigma\, A(\sigma) = o(1) < \tfrac{1}{2} \tag{50}$$

for sufficiently small ϵ. The failure probabilities in (49) and (50) sum to less than one, so some fixed choice of labels σ satisfies both requirements.

Finally, the outer OR adds one layer to the AND-rooted formula T. Since $\ln(st) = O_d(L)$,

$$\operatorname{depth}(C_m) \le d+1, \qquad |C_m| \le O(s|T|) = \exp(O_d(L)).$$

Simplify constant gates and flatten consecutive equal gates, retaining a unary AND gate when a child of the OR root reduces to a literal. This gives the asserted alternating presentation with AND-rooted children, without increasing the depth or asymptotic size. □

Now we are ready to prove Lemma 8.7.

Proof of Lemma 8.7. Take C_m from Lemma 8.12 and define $g = \neg C_m$. By De Morgan's laws, g is an AND-rooted circuit of depth at most $d+1$ and size $\exp(O_d(L))$. The monotonicity of C_m implies that g is antitone. For every X, either X or $\overline{X}$ has weight at least $m/2$, so the coverage property gives

$$C_m(X) \vee C_m(\overline{X}) = 1.$$

Consequently, $g(X) \wedge g(\overline{X}) = 0$, proving disjointness (42).

It remains to bound the dimension of U. For this gadget,

$$\mathcal{L} = \{S \subseteq [m] : C_m(1_S) = 0\}.$$

The space U is spanned by the multilinear monomials of degree at most k whose support in each block belongs to $\mathcal{L}$. We will bound the fraction of monomials excluded from U at each degree.

Put

$$\eta := \Pr_{X \sim \mathrm{Bern}(1/2-\epsilon)^m}[C_m(X) = 1].$$

We show that, for every integer $0 \le k \le (1/2-\epsilon)v$,

$$\gamma = \frac{\dim(U)}{H_v(k)} \ge 1 - (v+1)n\eta. \tag{51}$$

Fix $0 \le j \le k$ and choose a random support $S^{(j)}$ by including each of the v lifted variables independently with probability $p = j/v$. Write $S_i^{(j)}$ for its support in block i. Since $p \le 1/2 - \epsilon$ and C_m is monotone, each block satisfies

$$\Pr[S_i^{(j)} \notin \mathcal{L}] = \Pr_{X \sim \mathrm{Bern}(p)^m}[C_m(X) = 1] \le \eta.$$

A union bound over the n blocks therefore gives $\Pr[\exists i,\ S_i^{(j)} \notin \mathcal{L}] \le n\eta$.

Conditioned on $|S^{(j)}| = j$, the support $S^{(j)}$ is uniform among all subsets of size j. Moreover, $|S^{(j)}|$ has the binomial distribution with parameters v and j/v, for which j is a mode. Its probability at j is therefore at least $1/(v+1)$, since there are $v+1$ possible values. Thus the fraction of supports of size j that have a block outside $\mathcal{L}$ is at most

$$\Pr[\exists i,\ S_i^{(j)} \notin \mathcal{L} \mid |S^{(j)}| = j] \le \frac{\Pr[\exists i,\ S_i^{(j)} \notin \mathcal{L}]}{\Pr[|S^{(j)}| = j]} \le (v+1)n\eta.$$

There are $\binom{v}{j}$ multilinear monomials of degree j, and at least a fraction $1 - (v+1)n\eta$ of them belong to U for every $j \le k$. Summing over the degrees gives

$$\dim(U) \ge (1-(v+1)n\eta) \sum_{j=0}^{k} \binom{v}{j} = (1-(v+1)n\eta) H_v(k),$$

which proves (51). Finally, Lemma 8.12 gives $\eta \le \exp\left(-\epsilon^{-1/(d-1)}\right)$, proving the stated bound on γ. □

8.4 The size lower bound for lifted formulas

Proof of Theorem 8.1. We first choose the gadget so that it has polynomial size and the multiplier space occupies almost all of the available low-degree space. Apply Lemma 8.7 with

$$\epsilon := (32 \ln n)^{1-d}. \tag{52}$$

Take the block size m supplied by that lemma. The gadget then has size $n^{O_d(1)}$, and each block contains only $\Theta_d((\ln n)^{2d-1})$ variables. The dimension bound in the lemma will ensure that the multiplier space U has nearly the full dimension.

Since the base CNF has constant width, it has size $n^{O_{w_0}(1)}$. Thus Section 8.1 gives a lifted family of polynomial size whose axioms have depth at most d. Write

$$\begin{aligned} v &:= nm, \\ k &:= \lfloor (1/2 - \epsilon)v \rfloor. \end{aligned}$$

We next choose the approximation parameters to balance the size bound against the approximation degree. We use more repetitions at the top level to compensate for its larger error. For a sufficiently small constant $c_d > 0$, set

$$\begin{aligned} r &:= \left\lfloor c_d \ln n \, (\ln \Delta_n)^{1/(d-1)} \right\rfloor, \\ R &:= \lceil \sqrt{\Delta_n} \rceil. \end{aligned} \tag{53}$$

The choice of c_d ensures that

$$128 \epsilon r^{d-1} \le 128 (c_d/32)^{d-1} \ln \Delta_n \le \tfrac{1}{4} \ln \Delta_n$$

for sufficiently large n. This controls the loss at the top level. Meanwhile, r is polylogarithmic in n, so the approximation degree satisfies

$$q = R r^{d-1} = \sqrt{\Delta_n} \, (\ln n)^{O_d(1)} = o(\Delta_n).$$

Here the last equality uses $\Delta_n \ge n^C$: the polynomial growth of $\sqrt{\Delta_n}$ dominates the polylogarithmic factor. In particular, $3q + 1 < \Delta_n$, as required by the degree-transfer lemma.

The remaining parameter conditions in Corollary 4.18 hold for sufficiently large n: $\epsilon \le 1/256$, $v \ge 16/\epsilon^2$, and $r^{d-1} \le \epsilon v/2$. We apply the corollary with v in place of n. To estimate its exponent, note that $\epsilon r^{d-2} = O_d((\ln n)^{-1/(d-1)}) = o(1)$. Together with the bound on $128 \epsilon r^{d-1}$ above, this gives

$$\begin{aligned} \frac{1}{2} r e^{-128 \epsilon r^{d-2}} &\ge \frac{r}{4}, \\ \frac{1}{2} R e^{-128 \epsilon r^{d-1}} &\ge \frac{1}{2} \Delta_n^{1/4} \ge r. \end{aligned}$$

The last inequality holds for sufficiently large n, since $\Delta_n \ge n^C$ and r is polylogarithmic. Thus the exponent in the corollary is at least $r/4$. The gadget also leaves almost all candidate multipliers available: by Lemma 8.7,

$$\frac{\dim(U)}{H_v(k)} \ge 1 - (v+1) n^{-31} = 1 - o(1).$$

Finally, since $3q + 1 < \Delta_n$, Lemma 8.4 supplies the required non-zero relative pseudo-reduction operator for every non-zero multiplier in U. All hypotheses of Corollary 4.18 are therefore satisfied, and it gives

$$\mathrm{size}(\pi) \ge (1 - o(1)) \exp(r/4) \ge \exp\left(\frac{c_d}{8} \ln n \, (\ln \Delta_n)^{1/(d-1)} \right)$$

for every depth-d semantic $\mathsf{AC}^0[2]$-Frege refutation π and sufficiently large n, proving the claimed lower bound.

For the upper bound, successively conjoin the lifted axioms and then infer the constant-zero circuit. As in Section 8.1.1, every conjunction has depth at most $d + 1$, and sharing gives proof size linear in the total axiom size and number of axioms. All gates are AND or OR. □

Appendix

A Proof of the Pigeon-Degree Lower Bound

Theorem 3.7. *For all sufficiently large $n = 2^\ell$ and any $m > n$, any semantic polynomial calculus refutation of BPHP_n^m over $\mathbb{F}_2$ requires degree $n/6$.*

The result is a direct consequence of the unary pigeon-degree lower bound of Razborov [59]; see also [40]. We include the reduction because the statement for the binary encoding does not appear explicitly in these references.

We use the polynomial encoding of FPHP_n^m from Definition 6.6, with $n = 2^\ell$ and variables $x_{u,a}$ indexed by pigeons $u \in [m]$ and holes $a \in \{0,1\}^\ell$. For polynomials in these variables, pigeon degree is again the maximum number of different pigeons occurring in a monomial.

Theorem A.1 (Unary pigeon-degree lower bound [59, 40])**.** *If $m > n$, then FPHP_n^m has no PC refutations of pigeon degree at most $n/2$ over any field.*

We translate refutations of the bit pigeonhole principle into refutations of FPHP_n^m using the substitution

$$\tau(y_{u,t}) = \sum_{a : a_t = 1} x_{u,a}, \tag{54}$$

extended to polynomials by substitution.

Lemma A.2 (Decoding one pigeon)**.** *For every polynomial $q(y_{u,1}, \dots, y_{u,\ell})$ in the bits of one pigeon,*

$$\tau(q) \equiv \sum_{a \in \{0,1\}^\ell} q(a) x_{u,a}$$

modulo the pigeon, functionality, and Boolean axioms for pigeon u. The congruence has a PC derivation of pigeon degree one. In particular,

$$\tau(\delta_a(y_u)) \equiv x_{u,a}.$$

Proof. The pigeon and functionality axioms for pigeon u derive

$$\sum_{a \in \{0,1\}^\ell} x_{u,a} = 1$$

in pigeon degree one: expanding the pigeon axiom gives $1 - \sum_a x_{u,a}$ plus monomials divisible by functionality axioms, all involving only pigeon u.

Modulo the axioms for pigeon u, the variables $x_{u,a}$ are pairwise orthogonal idempotents whose sum is one. Moreover,

$$x_{u,a}\tau(y_{u,t}) \equiv a_t x_{u,a}.$$

Hence $x_{u,a}\tau(q) \equiv q(a)x_{u,a}$ for every a. Multiplying $\sum_a x_{u,a} \equiv 1$ by $\tau(q)$ and summing these identities gives the claim. All polynomials used in this calculation involve only pigeon u, so their pigeon degree is at most one. □

Lemma A.3 (Degree-preserving decoding)**.** *If BPHP_n^m has a PC refutation of pigeon degree at most d, then FPHP_n^m has a PC refutation of pigeon degree at most d.*

Proof. We translate the refutation line by line using τ. The image of a Boolean axiom for $y_{u,t}$ is a consequence of the axioms for pigeon u by Lemma A.2, and hence has a derivation of pigeon degree one.

Consider a collision axiom $\delta_a(y_u)\delta_a(y_{u'}) = 0$. Applying Lemma A.2 separately to pigeons u and u' gives

$$\tau\big(\delta_a(y_u)\delta_a(y_{u'})\big) \equiv x_{u,a}x_{u',a} = 0.$$

The last equality is a hole axiom of FPHP_n^m. This derivation has pigeon degree two.

Linear-combination inferences are preserved by linearity. Suppose that a source inference multiplies a polynomial p by $y_{u,t}$. From $\tau(p)$ we derive

$$\tau(y_{u,t})\tau(p) = \sum_{a:a_t=1} x_{u,a}\tau(p)$$

by multiplying by the variables $x_{u,a}$ and adding the results. Reducing using the axioms for pigeon u by Lemma A.2 yields the image of the multilinear polynomial $y_{u,t}p$. If a monomial of p already contains a variable of pigeon u, these steps do not enlarge its pigeon support. If it does not contain u, then its pigeon degree is at most $d-1$, since otherwise the multilinear source line $y_{u,t}p$ would have pigeon degree greater than d; cancellation of this monomial would require a monomial of p already containing the same d pigeons together with u. Thus every intermediate polynomial has pigeon degree at most d.

The final line 1 is mapped to 1, so the translated derivation is a refutation of FPHP_n^m of pigeon degree at most d. □

Proof of Theorem 3.7. Suppose that BPHP_n^m has a semantic PC refutation of degree at most $n/6$. By Lemma 2.1, it has a syntactic PC refutation of degree at most $3(n/6) = n/2$, and hence of pigeon degree at most $n/2$. By Lemma A.3, this gives a PC refutation of FPHP_n^m of pigeon degree at most $n/2$, contradicting Theorem A.1. □

B Proof of the Right-Degree Reduction Lemma

This appendix contains the proof of Lemma 6.3.

Proof of Lemma 6.3. Put $s = \gamma n$. For all sufficiently large n, choose an integer q such that

$$\frac{s}{64} \le q \le \frac{s}{32},$$

and let A_0 be the set of the q right vertices of largest degree.

Initially set $A = A_0$. Repeatedly choose a non-empty set J of at most $s/2$ remaining left vertices such that

$$|N_G(J) \setminus A| \le \frac{7}{4}|J|,$$

remove these vertices, and add their neighbours to A. If I is the set of left vertices removed so far, then

$$|N_G(I)| \le |A| \le q + \frac{7}{4}|I|.$$

Whenever $|I| \le s$, the neighbourhood bound gives

$$\frac{15}{8}|I| \le q + \frac{7}{4}|I|,$$

and hence $|I| \le 8q$. This bound holds throughout, since the next step leaves $|I| \le 8q + s/2 \le 3s/4$. Consequently,

$$|A| \le q + \frac{7}{4}|I| \le 15q \le \frac{15s}{32} < \frac{n}{2}.$$

We also have $|I| \le 8q \le s/4$. If I is non-empty, then

$$|A| \ge |N_G(I)| \ge \frac{15}{8}|I| > |I|;$$

if I is empty, the same strict inequality follows from $A_0 \subseteq A$.

By maximality of the cleaning process, every non-empty set J of at most $s/2$ left vertices of G' satisfies

$$|N_{G'}(J)| = |N_G(J) \setminus A| > \frac{7}{4}|J|.$$

Taking $|J| = 1$ shows that every left vertex of G' has degree two or three. For every non-empty set J of at most $s/2$ left vertices of G',

$$|\partial_{G'}(J)| \geq 2|N_{G'}(J)| - 3|J| > \frac{1}{2}|J|.$$

Finally, every right vertex outside A_0 has degree at most

$$\frac{3m}{q} \leq \frac{192\alpha}{\gamma}.$$

The claim follows by choosing an integer $d \geq 192\alpha/\gamma$. □

C Unsatisfiability and Expansion of Random Logarithmic-Width CNFs

We prove the first-moment estimate used in Lemma 7.1, keeping the dependence on the growing clause width explicit. Recall that $\ell = \lceil \log_2 n \rceil$, $m = n2^\ell$, and each of the m clauses is sampled independently and uniformly from the clauses of width exactly ℓ on distinct variables.

Proof of Lemma 7.1. The expected number of satisfying assignments is

$$2^n(1 - 2^{-\ell})^m \leq \exp(n \ln 2 - m2^{-\ell}) = \exp(-(1 - \ln 2)n) = o(1).$$

Hence F_n is unsatisfiable with probability $1 - o(1)$.

Let $s := \lfloor \eta n/\ell \rfloor$. Fix $1 \leq t \leq s$ and put $u = \lfloor 3\ell t/4 \rfloor$. If a set of t clauses has fewer than $3\ell t/4$ neighbours, all its variables lie in some set of u variables. For a fixed such set, a random clause uses only these variables with probability

$$\frac{\binom{u}{\ell}}{\binom{n}{\ell}} \leq \left(\frac{u}{n}\right)^\ell.$$

If $u < \ell$ the bad event is impossible. Otherwise the union bound gives

$$\Pr[\exists J :\ |J| = t,\ |N_{F_n}(J)| < 3\ell t/4] \leq \binom{m}{t}\binom{n}{u}\left(\frac{u}{n}\right)^{\ell t}. \tag{55}$$

Write $z = \ell t/n \leq \eta$. Using $\binom{a}{b} \leq (ea/b)^b$, the logarithm of the right-hand side, divided by t, is at most

$$1 + \ln \ell + \ell\left(\ln 2 + \frac{3}{4} + \frac{1}{4}\ln\frac{3}{4}\right) + \left(\frac{\ell}{4} - 1\right)\ln z. \tag{56}$$

To account for rounding u down, observe that $(en/u)^u(u/n)^{\ell t}$ increases with u for $0 < u < n$, so replacing u by $3\ell t/4$ only increases the bound. Choose $\eta > 0$ so small that

$$\ln 2 + \frac{3}{4} + \frac{1}{4}\ln\frac{3}{4} + \frac{1}{4}\ln\eta < -2.$$

For all sufficiently large n, (56) is at most $-\ell$, uniformly for $t \leq s$. Summing $\exp(-\ell t)$ over $1 \leq t \leq s$ proves the neighbourhood expansion with probability $1 - o(1)$.

Finally, since each clause has exactly ℓ neighbours,

$$|\partial_{F_n}(J)| \geq 2|N_{F_n}(J)| - \ell|J| \geq \frac{\ell}{2}|J|.$$

□

D Boundary Expanders for the Pigeonhole Principle

Proof of Lemma 7.4. We prove the claim by a direct application of Chernoff bounds and a union bound (cf. [61, proof of Theorem 2.5]). Fix a set I of $j \le s$ pigeons. For each hole, the probability of having exactly one neighbour in I is

$$\frac{j\Delta}{n}\left(1-\frac{\Delta}{n}\right)^{j-1} \ge \frac{3j\Delta}{4n}.$$

Indeed, since $j \le s \le n/(4\Delta)$, we have

$$\frac{j\Delta}{n} \le \frac{1}{4}.$$

Using Bernoulli's inequality, we obtain

$$\left(1-\frac{\Delta}{n}\right)^{j-1} \ge 1-\frac{(j-1)\Delta}{n} \ge 1-\frac{j\Delta}{n} \ge \frac{3}{4}.$$

Consequently,

$$\frac{j\Delta}{n}\left(1-\frac{\Delta}{n}\right)^{j-1} \ge \frac{j\Delta}{n}\cdot\frac{3}{4} = \frac{3j\Delta}{4n}.$$

We shall use the following standard Chernoff bounds. If X is a binomial random variable with mean μ, then

$$\Pr[X \le (1-\delta)\mu] \le \exp\left(-\frac{\delta^2\mu}{2}\right) \qquad (0 \le \delta \le 1),$$

and

$$\Pr[X \ge (1+\delta)\mu] \le \exp\left(-\frac{\delta^2\mu}{2+\delta}\right) \qquad (\delta \ge 0).$$

Since the events corresponding to distinct holes are independent, $X := |\partial_G(I)|$ is a binomial random variable. Its mean μ satisfies

$$\mu \ge \frac{3j\Delta}{4}.$$

In particular, $j\Delta/4 \le \mu/3$. Applying the lower-tail Chernoff bound with $\delta = 2/3$, we obtain

$$\Pr\left[X < \frac{j\Delta}{4}\right] \le \Pr\left[X \le \frac{\mu}{3}\right] \le \exp\left(-\frac{2\mu}{9}\right) \le \exp\left(-\frac{j\Delta}{6}\right).$$

Taking a union bound over all non-empty sets of at most s pigeons, the probability that the boundary-expansion condition fails is at most

$$\sum_{j=1}^{s}\binom{2n}{j}\exp(-j\Delta/6) \le \sum_{j=1}^{s}(2en\exp(-\Delta/6))^j = o(1).$$

The degree of a pigeon is binomial with mean Δ. Hence the upper-tail Chernoff bound with $\delta = 1$ gives

$$\Pr[\deg(u) > 2\Delta] \le \exp(-\Delta/3).$$

Similarly, the degree of a hole is binomial with mean 2Δ, and therefore

$$\Pr[\deg(v) > 4\Delta] \le \exp(-2\Delta/3).$$

A union bound over all $2n$ pigeons and n holes shows that the probability that any of these degree bounds fails is at most

$$2n\exp(-\Delta/3) + n\exp(-2\Delta/3) = o(1).$$

Together with the boundary-expansion estimate above, this proves that G has all the claimed properties asymptotically almost surely.

□

References

[1] Michael Alekhnovich. 'Lower Bounds for k-DNF Resolution on Random 3-CNFs'. In: *Computational Complexity* 20 (2011), pp. 597–614. DOI: 10.1007/s00037-011-0026-0.

[2] Michael Alekhnovich and Alexander A. Razborov. 'Lower Bounds for Polynomial Calculus: Non-Binomial Case'. In: *Proceedings of the Steklov Institute of Mathematics* 242 (2003), pp. 18–35.

[3] Yaroslav Alekseev and Nikita Gaevoy. *New Polynomial-Depth Res($\oplus$) Lower Bounds*. Electronic Colloquium on Computational Complexity, TR26-007. 2026.

[4] Yaroslav Alekseev, Svyatoslav Gryaznov, Toniann Pitassi, Robert Robere, Aleksandr Shekhovtsov and Iddo Tzameret. 'Progress on AC^0[2]-Frege Lower Bounds'. Manuscript. 2nd Nov. 2026.

[5] Yaroslav Alekseev and Dmitry Itsykson. 'Lifting to Bounded-Depth and Regular Resolutions over Parities via Games'. In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC 2025)*. ACM, 2025, pp. 584–595. DOI: 10.1145/3717823.3718150.

[6] Kazuyuki Amano. 'Bounds on the Size of Small Depth Circuits for Approximating Majority'. In: *Automata, Languages and Programming*. Vol. 5555. Lecture Notes in Computer Science. 2009, pp. 59–70. arXiv: 0902.0047.

[7] Albert Atserias, Maria Luisa Bonet and Juan Luis Esteban. 'Lower Bounds for the Weak Pigeonhole Principle and Random Formulas beyond Resolution'. In: *Information and Computation* 176 (2002), pp. 136–152. DOI: 10.1006/inco.2002.3114.

[8] Albert Atserias, Massimo Lauria and Jakob Nordström. 'Narrow Proofs May Be Maximally Long'. In: *ACM Transactions on Computational Logic* 17.3 (2016), 19:1–19:30. DOI: 10.1145/2898435.

[9] Albert Atserias and Moritz Müller. 'Automating Resolution is NP-Hard'. In: *Journal of the ACM* 67.5 (2020), 31:1–31:17. DOI: 10.1145/3409472.

[10] Paul Beame, Stephen A. Cook, Jeff Edmonds, Russell Impagliazzo and Toniann Pitassi. 'The Relative Complexity of NP Search Problems'. In: *Proceedings of the Twenty-Seventh Annual ACM Symposium on Theory of Computing*. ACM, 1995, pp. 303–314. DOI: 10.1145/225058.225147.

[11] Paul Beame, Russell Impagliazzo, Jan Krajíček, Toniann Pitassi and Pavel Pudlák. 'Lower Bounds on Hilbert's Nullstellensatz and Propositional Proofs'. In: *Proceedings of the London Mathematical Society* 73.1 (1996), pp. 1–26. DOI: 10.1112/plms/s3-73.1.1.

[12] Paul Beame and Sajin Koroth. 'On Disperser/Lifting Properties of the Index and Inner-Product Functions'. In: *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*. Vol. 251. Leibniz International Proceedings in Informatics. Schloss Dagstuhl, 2023, 14:1–14:17. DOI: 10.4230/LIPIcs.ITCS.2023.14.

[13] Eli Ben-Sasson, Prahladh Harsha and Sofya Raskhodnikova. 'Some 3CNF Properties Are Hard to Test'. In: *SIAM Journal on Computing* 35.1 (2005), pp. 1–21. DOI: 10.1137/S0097539704445445.

[14] Sreejata Kishor Bhattacharya, Farzan Byramji, Arkadev Chattopadhyay and Russell Impagliazzo. 'Lower Bounds for Near-Quadratic-Depth Resolution over Parities'. In: *Proceedings of the 58th Annual ACM Symposium on Theory of Computing (STOC 2026)*. ACM, 2026, pp. 955–966. DOI: 10.1145/3798129.3800809.

[15] Sreejata Kishor Bhattacharya, Arkadev Chattopadhyay and Pavel Dvořák. 'Exponential Separation Between Powers of Regular and General Resolution over Parities'. In: *39th Computational Complexity Conference (CCC 2024)*. Vol. 300. Leibniz International Proceedings in Informatics. Schloss Dagstuhl, 2024, 23:1–23:32. DOI: 10.4230/LIPIcs.CCC.2024.23.

[16] Kamil Braun. *An Exponential Lower Bound for the Bit Pigeonhole Principle in Resolution over Parities*. 2026. DOI: 10.48550/arXiv.2609.23015. arXiv: 2609.23015 [cs.CC].

[17] Samuel R. Buss. 'Lower Bounds on Nullstellensatz Proofs via Designs'. In: *Proof Complexity and Feasible Arithmetics*. Ed. by Paul W. Beame and Samuel R. Buss. Vol. 39. DIMACS Series in Discrete Mathematics and Theoretical Computer Science. American Mathematical Society, 1998, pp. 59–71. DOI: 10.1090/dimacs/039/04.

[18] Samuel R. Buss, Russell Impagliazzo, Jan Krajíček, Pavel Pudlák, Alexander A. Razborov and Jiří Sgall. 'Proof Complexity in Algebraic Systems and Bounded Depth Frege Systems with Modular Counting'. In: *Computational Complexity* 6.3 (1996), pp. 256–298. DOI: 10.1007/BF01294258.

[19] Samuel R. Buss, Leszek Aleksander Kołodziejczyk and Konrad Zdanowski. 'Collapsing Modular Counting in Bounded Arithmetic and Constant Depth Propositional Proofs'. In: *Transactions of the American Mathematical Society* 367.11 (2015), pp. 7517–7563. DOI: 10.1090/S0002-9947-2015-06233-3.

[20] Arkadev Chattopadhyay and Pavel Dvořák. 'Super-Critical Trade-Offs in Resolution over Parities via Lifting'. In: *40th Computational Complexity Conference (CCC 2025)*. Vol. 339. Leibniz International Proceedings in Informatics. Schloss Dagstuhl, 2025, 24:1–24:19. DOI: 10.4230/LIPIcs.CCC.2025.24.

[21] Arkadev Chattopadhyay, Nikhil S. Mande, Swagato Sanyal and Suhail Sherif. 'Lifting to Parity Decision Trees via Stifling'. In: *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*. Vol. 251. Leibniz International Proceedings in Informatics. Schloss Dagstuhl, 2023, 33:1–33:20. DOI: 10.4230/LIPIcs.ITCS.2023.33.

[22] Vašek Chvátal and Endre Szemerédi. 'Many Hard Examples for Resolution'. In: *Journal of the ACM* 35.4 (1988), pp. 759–768. DOI: 10.1145/48014.48016.

[23] Matthew Clegg, Jeff Edmonds and Russell Impagliazzo. 'Using the Gröbner Basis Algorithm to Find Proofs of Unsatisfiability'. In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. ACM, 1996, pp. 174–183. DOI: 10.1145/237814.237860.

[24] Jonas Conneryd. 'Random $(\log n)$-CNF Formulas Are Hard for Depth-2 Frege with Modular Gates'. In: (Oct. 2026). DOI: 10.48550/arXiv.2610.40397. arXiv: 2610.40397 [cs.CC].

[25] Stephen A. Cook and Robert A. Reckhow. 'The Relative Efficiency of Propositional Proof Systems'. In: *Journal of Symbolic Logic* 44.1 (1979), pp. 36–50. DOI: 10.2307/2273702.

[26] Klim Efremenko, Michal Garlík and Dmitry Itsykson. 'Lower Bounds for Regular Resolution over Parities'. In: *SIAM Journal on Computing* 54.4 (2025), pp. 887–915. DOI: 10.1137/24M1696640.

[27] Klim Efremenko and Dmitry Itsykson. 'Amortized Closure and Its Applications in Lifting for Resolution over Parities'. In: *40th Computational Complexity Conference (CCC 2025)*. Vol. 339. Leibniz International Proceedings in Informatics. Schloss Dagstuhl, 2025, 8:1–8:24. DOI: 10.4230/LIPIcs.CCC.2025.8.

[28] Klim Efremenko and Dmitry Itsykson. 'Strong ETH Holds for Bounded-Depth Resolution over Parities'. In: *Proceedings of the 58th Annual ACM Symposium on Theory of Computing (STOC 2026)*. ACM, 2026, pp. 898–909. DOI: 10.1145/3798129.3800804.

[29] Yuval Filmus. 'Another Look at Degree Lower Bounds for Polynomial Calculus'. In: *Theoretical Computer Science* 796 (2019), pp. 286–293. DOI: 10.1016/j.tcs.2019.09.023.

[30] John Franco and Marvin C. Paull. 'Probabilistic analysis of the Davis Putnam procedure for solving the satisfiability problem'. In: *Discret. Appl. Math.* 5.1 (1983), pp. 77–87. DOI: 10.1016/0166-218X(83)90017-3.

[31] Ehud Friedgut. 'Sharp thresholds of graph properties, and the k-sat problem'. In: *J. Amer. Math. Soc.* 12.4 (1999). With an appendix by Jean Bourgain, pp. 1017–1054. DOI: 10.1090/S0894-0347-99-00305-7.

[32] Michal Garlík and Leszek Aleksander Kołodziejczyk. 'Some Subsystems of Constant-Depth Frege with Parity'. In: *ACM Transactions on Computational Logic* 19.4 (2018), 29:1–29:34. DOI: 10.1145/3243126.

[33] Mika Göös, Sajin Koroth, Ian Mertz and Toniann Pitassi. 'Automating Cutting Planes is NP-Hard'. In: *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2020. Association for Computing Machinery, 2020, pp. 68–77. DOI: 10.1145/3357713.3384248.

[34] Dima Grigoriev and Edward A. Hirsch. 'Algebraic Proof Systems over Formulas'. In: *Theoretical Computer Science* 303.1 (2003), pp. 83–102. DOI: 10.1016/S0304-3975(02)00446-2.

[35] Joshua A. Grochow and Toniann Pitassi. 'Circuit Complexity, Proof Complexity, and Polynomial Identity Testing: The Ideal Proof System'. In: *Journal of the ACM* 65.6, 37 (2018), pp. 1–59. DOI: 10.1145/3230742.

[36] Svyatoslav Gryaznov. 'Notes on Resolution over Linear Equations'. In: *Computer Science – Theory and Applications (CSR 2019)*. Vol. 11532. Lecture Notes in Computer Science. Springer, 2019, pp. 168–179. DOI: 10.1007/978-3-030-19955-5_15.

[37] Svyatoslav Gryaznov, Sergei Ovcharov and Artur Riazanov. 'Resolution over Linear Equations: Combinatorial Games for Tree-like Size and Space'. In: *ACM Transactions on Computation Theory* 16.3 (2024), 15:1–15:15. DOI: 10.1145/3675415.

[38] Tuomas Hakoniemi, Nutan Limaye and Iddo Tzameret. 'Functional Lower Bounds in Algebraic Proofs: Symmetry, Lifting, and Barriers'. In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. STOC 2024. Association for Computing Machinery, 2024, pp. 1396–1404. DOI: 10.1145/3618260.3649616.

[39] Russell Impagliazzo, Sasank Mouli and Toniann Pitassi. 'The Surprising Power of Constant Depth Algebraic Proofs'. In: *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science*. ACM, 2020, pp. 591–603. DOI: 10.1145/3373718.3394754.

[40] Russell Impagliazzo, Pavel Pudlák and Jiří Sgall. 'Lower Bounds for the Polynomial Calculus and the Gröbner Basis Algorithm'. In: *Computational Complexity* 8 (1999), pp. 127–144. DOI: 10.1007/s000370050024.

[41] Dmitry Itsykson and Alexander Knop. 'Supercritical Tradeoff Between Size and Depth for Resolution over Parities'. In: *17th Innovations in Theoretical Computer Science Conference (ITCS 2026)*. Vol. 362. Leibniz International Proceedings in Informatics. Schloss Dagstuhl, 2026, 81:1–81:20. DOI: 10.4230/LIPIcs.ITCS.2026.81.

[42] Dmitry Itsykson, Vladimir Podolskii and Alexander Shekhovtsov. 'Resolution Width Lifts to Near-Quadratic-Depth Res(⊕) Size'. In: *41st Computational Complexity Conference (CCC 2026)*. Vol. 383. Leibniz International Proceedings in Informatics. Schloss Dagstuhl, 2026, 13:1–13:27. DOI: 10.4230/LIPIcs.CCC.2026.13.

[43] Dmitry Itsykson and Artur Riazanov. 'Proof Complexity of Natural Formulas via Communication Arguments'. In: *36th Computational Complexity Conference (CCC 2021)*. Vol. 200. Leibniz International Proceedings in Informatics. Schloss Dagstuhl, 2021, 3:1–3:34. DOI: 10.4230/LIPIcs.CCC.2021.3.

[44] Dmitry Itsykson and Dmitry Sokolov. 'Lower Bounds for Splittings by Linear Combinations'. In: *Mathematical Foundations of Computer Science 2014*. Vol. 8635. Lecture Notes in Computer Science. Springer, 2014, pp. 372–383. DOI: 10.1007/978-3-662-44465-8_32.

[45] Dmitry Itsykson and Dmitry Sokolov. 'Resolution over Linear Equations Modulo Two'. In: *Annals of Pure and Applied Logic* 171.1, 102722 (2020). DOI: 10.1016/j.apal.2019.102722.

[46] Ondřej Ježil and Dimitrios Tsintsilidas. Personal communication. 2026.

[47] Erfan Khaniki. 'On Proof Complexity of Resolution over Polynomial Calculus'. In: *ACM Trans. Comput. Log.* 23.3 (2022), 16:1–16:24. DOI: 10.1145/3506702.

[48] Jan Krajíček. *Bounded Arithmetic, Propositional Logic, and Complexity Theory*. Vol. 60. Encyclopedia of Mathematics and Its Applications. Cambridge University Press, 1995. DOI: 10.1017/CBO9780511529948.

[49] Jan Krajíček. *Proof Complexity*. Vol. 170. Encyclopedia of Mathematics and Its Applications. Cambridge University Press, 2019. DOI: 10.1017/9781108242066.

[50] Jan Krajíček. 'Randomized Feasible Interpolation and Monotone Circuits with a Local Oracle'. In: *Journal of Mathematical Logic* 18.2, 1850012 (2018). DOI: 10.1142/S0219061318500125.

[51] Massimo Lauria and Jakob Nordström. 'Tight Size-Degree Bounds for Sums-of-Squares Proofs'. In: *Computational Complexity* 26.4 (2017), pp. 911–948. DOI: 10.1007/s00037-017-0152-4.

[52] Nutan Limaye, Karteek Sreenivasaiah, Srikanth Srinivasan, Utkarsh Tripathi and S. Venkitesh. 'A Fixed-Depth Size-Hierarchy Theorem for $\mathsf{AC}^0[\oplus]$ via the Coin Problem'. In: *SIAM Journal on Computing* 50 (2021). DOI: 10.1137/19M1276467.

[53] Alexis Maciel, Toniann Pitassi and Alan R. Woods. 'A New Proof of the Weak Pigeonhole Principle'. In: *Journal of Computer and System Sciences* 64.4 (2002), pp. 843–872. DOI: 10.1006/jcss.2002.1830.

[54] Mladen Mikša and Jakob Nordström. 'A Generalized Method for Proving Polynomial Calculus Degree Lower Bounds'. In: *Journal of the ACM* 71.6, 37 (2024), pp. 1–43. DOI: 10.1145/3675668.

[55] Shuo Pang. *A Degree–Size Relation for Resolution over Polynomials*. 2026. arXiv: 2610.00837 [cs.CC].

[56] Toniann Pitassi. 'Algebraic Propositional Proof Systems'. In: *Descriptive Complexity and Finite Models*. Ed. by Neil Immerman and Phokion G. Kolaitis. Vol. 31. DIMACS Series in Discrete Mathematics and Theoretical Computer Science. American Mathematical Society, 1997, pp. 215–244. DOI: 10.1090/dimacs/031/07.

[57] Ran Raz and Iddo Tzameret. 'Resolution over Linear Equations and Multilinear Proofs'. In: *Annals of Pure and Applied Logic* 155.3 (2008), pp. 194–224. DOI: 10.1016/j.apal.2008.04.001.

[58] Ran Raz and Iddo Tzameret. 'The Strength of Multilinear Proofs'. In: *Computational Complexity* 17.3 (2008), pp. 407–457. DOI: 10.1007/s00037-008-0246-0.

[59] Alexander A. Razborov. 'Lower Bounds for the Polynomial Calculus'. In: *Computational Complexity* 7 (1998), pp. 291–324. DOI: 10.1007/s000370050013.

[60] Alexander A. Razborov. 'Lower Bounds on the Size of Bounded Depth Circuits over a Complete Basis with Logical Addition'. In: *Mathematical Notes* 41.4 (1987), pp. 333–338. DOI: 10.1007/BF01137685.

[61] Alexander A. Razborov. 'Pseudorandom Generators Hard for k-DNF Resolution and Polynomial Calculus Resolution'. In: *Annals of Mathematics* 181.2 (2015), pp. 415–472. DOI: 10.4007/annals.2015.181.2.1.

[62] Susanna F. de Rezende, David Engström, Yassine Ghannane, Duri Andrea Janett and Artur Riazanov. 'Average-Case Hardness of Binary-Encoded Clique in Proof and Communication Complexity'. In: *53rd International Colloquium on Automata, Languages, and Programming (ICALP 2026)*. Vol. 374. Leibniz International Proceedings in Informatics. Schloss Dagstuhl, 2026, 151:1–151:23. DOI: 10.4230/LIPIcs.ICALP.2026.151.

[63] Susanna F. de Rezende, David Engström, Yassine Ghannane and Kilian Risse. *Superpolynomial Length Lower Bounds for Tree-Like Semantic Proof Systems with Bounded Line Size*. Electronic Colloquium on Computational Complexity, TR26-078. 2026.

[64] Roman Smolensky. 'Algebraic Methods in the Theory of Lower Bounds for Boolean Circuit Complexity'. In: *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*. ACM, 1987, pp. 77–82. DOI: 10.1145/28395.28404.

[65] Iddo Tzameret. 'Algebraic Proofs over Noncommutative Formulas'. In: *Information and Computation* 209.10 (2011), pp. 1269–1292. DOI: 10.1016/j.ic.2011.07.004.