Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 17 (2024)

On Pigeonhole Principles and Ramsey in TFNP

Siddhartha Jain
UT Austin

Jiawei Li
UT Austin

Robert Robere
McGill

Zhiyang Xun
UT Austin

August 8, 2024

Abstract

We show that the TFNP problem Ramsey is not black-box reducible to Pigeon, refuting a conjecture of Goldberg and Papadimitriou in the black-box setting. We prove this by giving reductions to Ramsey from a new family of TFNP problems that correspond to generalized versions of the pigeonhole principle, and then proving that these generalized versions cannot be reduced to Pigeon. Formally, we define t-PPP as the class of total NP-search problems reducible to finding a t-collision in a mapping from $(t-1)N+1$ pigeons to N holes. These classes are closely related to multi-collision resistant hash functions in cryptography. We show that the generalized pigeonhole classes form a hierarchy as t increases, and also give a natural condition on the parameters t_1, t_2 that captures exactly when t_1-PPP and t_2-PPP collapse in the black-box setting. Finally, we prove other inclusion and separation results between these generalized Pigeon problems and other previously studied TFNP subclasses, such as PLS, PPA, and PLC. Our separation results rely on new lower bounds in propositional proof complexity based on pseudoexpectation operators, which may be of independent interest.

1 Introduction

The theory of TFNP is the study of NP search problems that are *guaranteed* to have solutions. In most problems studied in the literature, this guarantee is usually due to a non-constructive combinatorial lemma [JPY88, Pap94]. Perhaps the most famous example is the Pigeonhole Principle (PHP), which defines the search problem Pigeon: given a polynomial-size circuit encoding a mapping from $N+1$ pigeons to N holes, find two pigeons that are in the same hole. However, many other non-constructive combinatorial principles also play important roles in TFNP, including variants of the Handshaking Lemma (corresponding to the classes PPA, PPAD) [Pap94, BCE+98], the guaranteed convergence of local search algorithms (corresponding to the class PLS) [JPY88, Pap94], and the guaranteed convergence of gradient descent (corresponding to the class CLS [DP11, FGHS23]). Importantly, *polynomial-time reducibility* between these various search problems corresponds directly to how *relatively constructive* the corresponding combinatorial principles are. From this viewpoint, TFNP can be seen as a kind of "bounded reverse mathematics", in which we seek to study what families of non-constructive combinatorial principles can have their witnesses *constructively* reduced to each other.

One problem which has so far resisted classification is the search problem corresponding to *Ramsey's Theorem*. In the Ramsey problem, introduced by Krajíček [Kra05], we are given a polynomial-size circuit C encoding the edge relation of a graph on $N = 2^{n+1}$ vertices, and our goal

ISSN 1433-8092

is to find a clique or independent set of size n on the graph. Since the underlying combinatorial principle comes from extremal combinatorics, it is natural to ask if there is a polynomial-time reduction to or from the Pigeonhole Principle, the prototypical example of extremal reasoning. This seems especially reasonable given that standard proof of Ramsey's theorem is essentially a *recursive* application of the pigeonhole principle[1]. Whether or not RAMSEY is in PPP— the TFNP subclass defined by polynomial-time reductions PIGEON— has been informally asked many times, was formally conjectured by Goldberg and Papadimitriou [GP17], and also appeared in Daskalakis's recent ICM plenary lecture [Das19, Open Question 17].

Conjecture 1 (Goldberg & Papadimitriou [GP17])**.** RAMSEY is in PPP.

However, despite the attention given to this problem, actually finding a reduction from RAMSEY to PPP has remained elusive. Recently, a line of work has weakened this goal, and instead attempted to place RAMSEY into *any* natural subclass of TFNP [KoT22, PPY23, BFH+23]. Of particular note is the work of Pasarkar, Papadimitriou, and Yannakakis [PPY23], who defined a novel new TFNP subclass called PLC. This class exactly captures the kinds of "recursive applications" of the pigeonhole principle that are used in the typical proof of Ramsey's Theorem and other extremal combinatorial principles, like the Erdös-Rado Sunflower Lemma. Pasarkar et al. showed that both RAMSEY and PIGEON were contained in PLC, but left open the question of determining how much extra power PLC contains over PPP.

1.1 Our Results

Our main result is to negatively resolve Conjecture 1 in the *black-box setting*, where the inputs are provided by oracle queries.

Theorem 1.1. There is no black-box reduction from RAMSEY to PIGEON.

We note that *all* known upper-bound techniques in TFNP also work in the black-box setting, so our theorem rules out any approach to the above conjecture using current technology. Any unconditional white-box separation within TFNP would imply $\mathsf{P} \neq \mathsf{NP}$, thus, one can only hope to prove separations in the black-box setting without a breakthrough in complexity theory. Also, as an immediate corollary of our previous theorem, we conclude that PLC strictly contains PPP with respect to black-box reductions, answering a question of Pasarkar, Papadimitriou, and Yannakakis [PPY23].

To prove Theorem 1.1, we introduce a new family of problems in TFNP that correspond to *generalized* Pigeonhole Principles and systematically study their properties. Whereas the Pigeonhole Principle says that any method of placing $n+1$ pigeons into n holes must result in a collision of *two* pigeons in one hole, the generalized Pigeonhole Principle we study says that any method of placing $t(n+1)$ pigeons into n holes must result in a collision of $t+1$ pigeons in a hole. We can encode this as a total search problem as follows.

Definition 1.2. For any positive integer n let $t(n), M(n)$, and $N(n)$ be integer parameters satisfying $M > (t-1)N$. The t-PIGEON$_N^M$ problem is defined as follows.

> **Input** (n, h), where n is given in unary and $h : [M] \to [N]$ is a map of M pigeons to N holes represented by a poly(n)-size circuit.

[1] Start by picking a vertex v, and then delete all vertices that are adjacent to v or non-adjacent to v, whichever set is larger, and then repeat this process. Continuing in this way we can construct a sequence of $\log n$ vertices, and at least half of these must be either all adjacent or non-adjacent to each other.

Solutions A *t-collision* in h, which is a set of t pigeons that are all assigned to the same hole by the circuit.

We note that for this problem to be in TFNP, $t(n)$ must have at most polynomial growth rate, as otherwise a solution will be too large to verify in polynomial time. As an important special case, we use t-PIGEON to indicate t-PIGEON$_N^M$ with the tight parameter setting $M = (t-1)N+1$ — note that 2-PIGEON = PIGEON, for instance. Finally, $N = 2^n$ in the typical setting of interest, and so we will use "n" and "$\log N$" interchangeably as parameters.

The PIGEON problem and its t-PIGEON variants are naturally related to cryptography, as solving them efficiently would imply the ability to find collisions in collision-resistant hash functions. These hash functions usually compress the output from say $2n$ to n bits, hence ore relevant here is the 'weak' version of PIGEON, which defines the class PWPP. While they have been implicitly studied before by cryptographers (see Related Work in Section 1.3 for more details), they are studied systematically[2] here from the perspective of TFNP for the first time. To relate these problems to RAMSEY, we generalize an argument due to Komargodski, Naor, and Yogev [KNY18] to obtain the following:

Theorem 1.3. Whenever $M \geq N^{4t}/4^t$, t-PIGEON$_N^M$ can be black-box reduced to RAMSEY and its bipartite variant BIRAMSEY.

We note that Krajíček [Kra05] showed that 2-PIGEON$_N^{2N}$ reduces to RAMSEY, which is generalized by the above theorem. The previous theorem implies, for instance, that there is a black-box reduction from t-PIGEON$_N^{N^{4t}}$ to RAMSEY, and so to prove Theorem 1.1 we will argue that there is *no* black-box reduction from t-PIGEON$_N^{N^{4t}}$ to PIGEON. In order to do this we develop and apply techniques from *propositional proof complexity*, which are described in more detail in Section 1.2 and may be of independent interest. For now, we note that t-PIGEON$_N^{N^{4t}}$ is, at first glance, seemingly incomparable with PIGEON in power. On one hand, t-PIGEON$_N^{N^{4t}}$ is much *weaker* than PIGEON, as the number of pigeons is so much greater than the number of holes. On the other hand, in this problem we are seeking a collision of t pigeons in a hole, and it is not at all clear how to do this when we can only guarantee a collision of *two* pigeons.

1.1.1 The Pecking Order

Indeed, we will prove much more than the non-reducibility of t-PIGEON$_N^{N^{4t}}$ to PIGEON. In fact, we exhibit a *hierarchy theorem*: t-PIGEON$_N^M$ does not black-box reduce to $(t-1)$-PIGEON$_{N'}^{M'}$, *no matter the compression rates of pigeons to holes in either problem.* Before we formally state our next results, let us first introduce complexity classes capturing these generalized pigeon problems.

Definition 1.4 (***t*-PPP and *t*-PWPP)**. For any function $t(n) \geq 2$, define the classes

***t*-PPP.** All search problems reducible to t-PIGEON.

***t*-PWPP.** All search problems reducible to t-PIGEON$_N^M$, with $M = (t-1+c)N$ for constant $c > 0$.

Note that we can define the classic TFNP classes PPP and PWPP by taking $t = 2$ in Definition 1.4. Trivially, we have t-PWPP $\subseteq$ t-PPP for any $t \geq 2$. Furthermore, a simple padding argument (cf. Lemma 2.12) shows that these classes form a hierarchy: t-PPP is contained in $(t+1)$-PPP for each t. We call the hierarchy of classes for *constant* values of t the *Pigeon Hierarchy*, denoted by PiH.

[2] Also see the recent independent work of [BGS24], in which these problems are also studied.

Definition 1.5 (Pigeon Hierarchy)**.** $\mathsf{PiH} = \bigcup_{t=2}^{\infty} t\text{-}\mathsf{PPP}$

Lying atop the Pigeon Hierarchy are the t-PPP classes where $t(n)$ is a *growing* function of the input size. We isolate two interesting cases here. The strongest we call PAP, for *Polynomial Averaging Principle*.

Definition 1.6. PAP is the set of all search problems reducible to n-PPP.

By a fairly simple argument (cf. Theorem 1.10), one can show that $t(n)$-PPP is contained in PAP for *any* polynomial function $t(n)$. Since $t(n)$ must have polynomial growth rate in order for t-PIGEON to be in TFNP, this means that PAP is the strongest possible class of generalized pigeon problems. Similar to this definition, we introduce the class SAP (for *Subpolynomial Averaging Principle*).

Definition 1.7. SAP is the set of all search problems reducible to $t(n)$-PPP, for some $t(n)$ sub-polynomial[3] in n.

Intuitively, SAP contains all problems defined by t-PIGEON, but excludes the hardest problems in PAP, which is convenient when stating our results in the strongest form. We call this entire collection of complexity classes the *Pecking Order*.

1.1.2 Structure of the Pecking Order

We are able to *completely* characterize the relationships between various t-PPP classes in the black-box setting. In the black-box setting, instead of providing the inputs succinctly via polynomial-size boolean circuits, we instead think of the inputs as being provided as black-box oracles. For example, in the black-box version of the PIGEON problem, the map from pigeons to holes is provided by an oracle f which, given the name of the pigeon, outputs the hole that the pigeon maps to. In general, if A is a total search problem defining a TFNP subclass A via black-box reducibility, we will use A^{dt} (A^{dt}, resp.) to denote the black-box versions of this problem and class (we refer to Section 2 for formal definitions).

To start, on the side of separations, we are able to prove that the Pigeon Hierarchy (PiH) is strict. Indeed, we can prove very strong black-box separations between t-$\textsc{Pigeon}_N^M$ and $(t+1)$-$\textsc{Pigeon}_{N'}^{M'}$.

Theorem 1.8. If t is a constant and M, N, M', N' are parameters chosen so that $M' = \mathrm{poly}(M)$, $M \geq tN + 1$, $M' \geq (t-1)N' + 1$, then $(t+1)$-$\textsc{Pigeon}_N^M$ does not have an efficient black-box reduction to t-$\textsc{Pigeon}_{N'}^{M'}$. In particular, $(t+1)\text{-}\mathsf{PPP}^{dt} \not\subseteq t\text{-}\mathsf{PPP}^{dt}$ for any constant t, and so PiH^{dt} forms a strict hierarchy in the black-box setting.

Note that the above separation is invariant of compression rate, and so it even shows the stronger separation $(t+1)\text{-}\mathsf{PWPP}^{dt} \not\subseteq t\text{-}\mathsf{PPP}^{dt}$. We can also prove separation results for non-constant (growing) t. Before stating this generalization, let us introduce a helpful definition for relating the various collision rates.

Definition 1.9. A function $b(n)$ is *polynomially close* to $a(n)$ if there exists a polynomial $p(n)$ such that for any n, $b(n) \leq a(p(n))$, and $a(n) \leq b(p(n))$.

For example, all functions with polynomial growth rate are polynomially close to each other, while slower-growing functions like poly-logarithms are not. This definition is useful since it captures the t-PPP classes in the following sense:

[3] A function $t(n)$ is said to be sub-polynomial if $t(n) = o(n^c)$ for any $c > 0$.

Theorem 1.10. If a function $b(n)$ is polynomially close to $a(n)$, then a-PPP $=$ b-PPP.

The proof of Theorem 1.10 uses a fairly standard *copying argument* to manipulate the parameters, and we leave the details to Section 2.3. By generalizing our separation result for the Pigeon Hierarchy (Theorem 1.8), we can prove a converse to Theorem 1.10.

Theorem 1.11. Let $a(n), b(n)$ be functions such that $a(n)$ is larger than $b(n)$ and is not polynomially close to $b(n)$. Let M, N, M', N' be any parameters chosen so that $M' = \text{poly}(M)$, $M \geq (a(n)-1)N+1$, $M' \geq (b(n)-1)N'+1$. Then $a(n)$-PIGEON_N^M does not have an efficient black-box reduction to $b(n)$-$\text{PIGEON}_{N'}^{M'}$.

Since the constant function $a(n) = (t+1)$ is not polynomially close to $b(n) = t$ for any constant t, this is a strict generalization of Theorem 1.8. Combining Theorem 1.10 with Theorem 1.11, we can completely characterize the structure of the Pecking Order with respect to the collision number.

Theorem 1.12. For any two functions $a(n), b(n)$, $a(n)$-$\text{PPP}^{dt} = b(n)$-PPP^{dt} if and only if $a(n)$ and $b(n)$ are polynomially close.

We emphasize that Theorem 1.8 and Theorem 1.11 do not depend on the specific choice of parameters M, N, M', N', as long as they are non-trivial. For example, 3-$\text{PIGEON}_N^{N^2}$ cannot be reduced to PIGEON_N^{N+1}, though the latter one has a much smaller compression rate. Conceptually:

We cannot trade a lower compression rate in exchange for more collisions.

This favorable property of our structural theorem makes it very convenient for showing separation results for problems from the Pecking Order. Once we can reduce t-PIGEON_N^M to a search problem of interest (like RAMSEY) for *any* value of $M \gg t \cdot N$, we *automatically* separate the search problem from any lower level of the Pecking Order than t-PPP, including PPP.

Conversely, we show that one cannot perform a tradeoff in the *other* direction (even in the randomized setting). Formally speaking, we separate PPP from n-PWPP for any randomized black-box reduction.

Theorem 1.13. There is no randomized black-box reduction from PIGEON to n-PIGEON_N^M with $M \geq (n-1+c)N$ for a constant $c > 0$.

Complementary to Theorem 1.11, this result can be interpreted as saying:

We cannot trade more collisions in exchange for an (extremely) low compression rate.

In a concurrent work by Li [Li24], Theorem 1.13 is generalized to total search problems with many solutions on any given input, like PWPP and variants. We refer to Section 1.3 for further related discussion.

We finally remark that our separations between the Pecking Order classes can be viewed as a first step towards a notorious problem in cryptography: determining whether there is a *black-box construction* of *multi-collision resistant hash-functions* from *collision-resistant hash-functions* (cf. Section 1.3). In our language, this (very roughly) translates to proving or disproving the existence of a *randomized* black-box Turing reductions from t-PIGEON_N^M to PIGEON. Our main separation result can be seen as a confirmation that there is no *deterministic* black-box reduction, and Theorem 1.13 is a weak (many-one) separation in the converse direction. However, we do not obtain a randomized separation in this paper, and leave it open for future work.

1.1.3 Ramsey and the Pecking Order

By employing our separations for Pecking Order classes, we give strong lower bounds for the Ramsey problem in the black-box setting, as well as for its bipartite variant BiRamsey, where we are given a bipartite graph on (N, N) vertices, and have to output either a $(\log N)/2$-biclique or a $(\log N)/2$-bi-independent set. In particular, combining Theorem 1.3 with Theorem 1.11 we obtain the following black-box separation, which is significantly stronger than Theorem 1.1.

Theorem 1.14. $\text{Ramsey}^{dt}, \text{BiRamsey}^{dt} \notin \mathsf{SAP}^{dt}$. In particular, $\text{Ramsey}^{dt}, \text{BiRamsey}^{dt} \notin \mathsf{PPP}^{dt}$.

In the reverse direction, we show that BiRamsey fits into the Pecking Order with a slightly weaker parameter. Together with Theorem 1.14, we give an almost tight characterization of BiRamsey using the Pecking Order.

Theorem 1.15. (Generalization of [KNY19, Theorem 3.10]) $\frac{n-\log n}{2}$-$\text{BiRamsey} \in \mathsf{PAP}$.

We also reveal an interesting connection between the Pecking Order and the *Polynomial Long Choice* class (PLC) introduced by [PPY23]. Recall that [PPY23] has shown that $\text{Ramsey} \in \mathsf{PLC}$, therefore, by Theorem 1.14, we already separate PLC from PPP in the black-box setting, resolving one of their open questions. However, we can get even stronger results by considering a subclass of PLC called UPLC [PPY23] (Unary PLC, see Definition 6.14), which is the *non-adaptive* variant of PLC. In particular, we can show that $n/2$-$\text{Pigeon}^{N}_{\sqrt{N}}$ reduces to UPLC (Lemma 6.15), and therefore we can separate UPLC from SAP (and of course, PPP) in the black-box setting using our separation results.

Theorem 1.16. $\mathsf{UPLC}^{dt} \not\subseteq \mathsf{SAP}^{dt}$.

Conversely, we show in Lemma 6.16 that $\mathsf{UPLC} \subseteq n$-PWPP, and thus by our separation regarding to the compression rate (Theorem 1.13), we have that $\mathsf{PPP}^{dt} \not\subseteq \mathsf{UPLC}^{dt}$, resolving another open question of Pasarkar, Papadimitriou, and Yannakakis.

Theorem 1.17. $\mathsf{PPP}^{dt} \not\subseteq \mathsf{UPLC}^{dt}$. Consequently, $\mathsf{PLC}^{dt} \not\subseteq \mathsf{UPLC}^{dt}$.

We also observe that UPLC does not capture the full strength of the non-adaptive iterative pigeonhole principle. We define a new problem T-UPLC (Definition 6.17), which is similar to UPLC, but has the strongest possible parameters. We show that T-UPLC is indeed PAP-complete (Lemma 6.18), therefore, the non-adaptive iterative pigeonhole principle is equivalent to the generalized pigeonhole principle.

Taken together, our results provide a nearly complete picture (see Figure 1 for a summary) of the relative complexities of the Pecking Order, PLC, UPLC, Ramsey, and BiRamsey in the black-box setting.

1.1.4 PLS, PPA, and the Pecking Order

Given these results, it is also quite natural to try and relate the Pecking Order classes to other well-studied and important TFNP subclasses. We focus on the relationships with the classes PLS and PPA, which are the two of the three strongest TFNP subclasses out of the "original five" TFNP subclasses defined by Papadimitriou [Pap94] (the third being PPP). The seminal work of Beame et. al. [BCE+98] has shown that the important TFNP subclass PPA, embodying the *Handshaking Lemma*, is not contained in PPP in the black-box setting. The recent breakthrough of [GHJ+22] showed that the TFNP subclass PLS, embodying problems with (not necessarily efficient) local

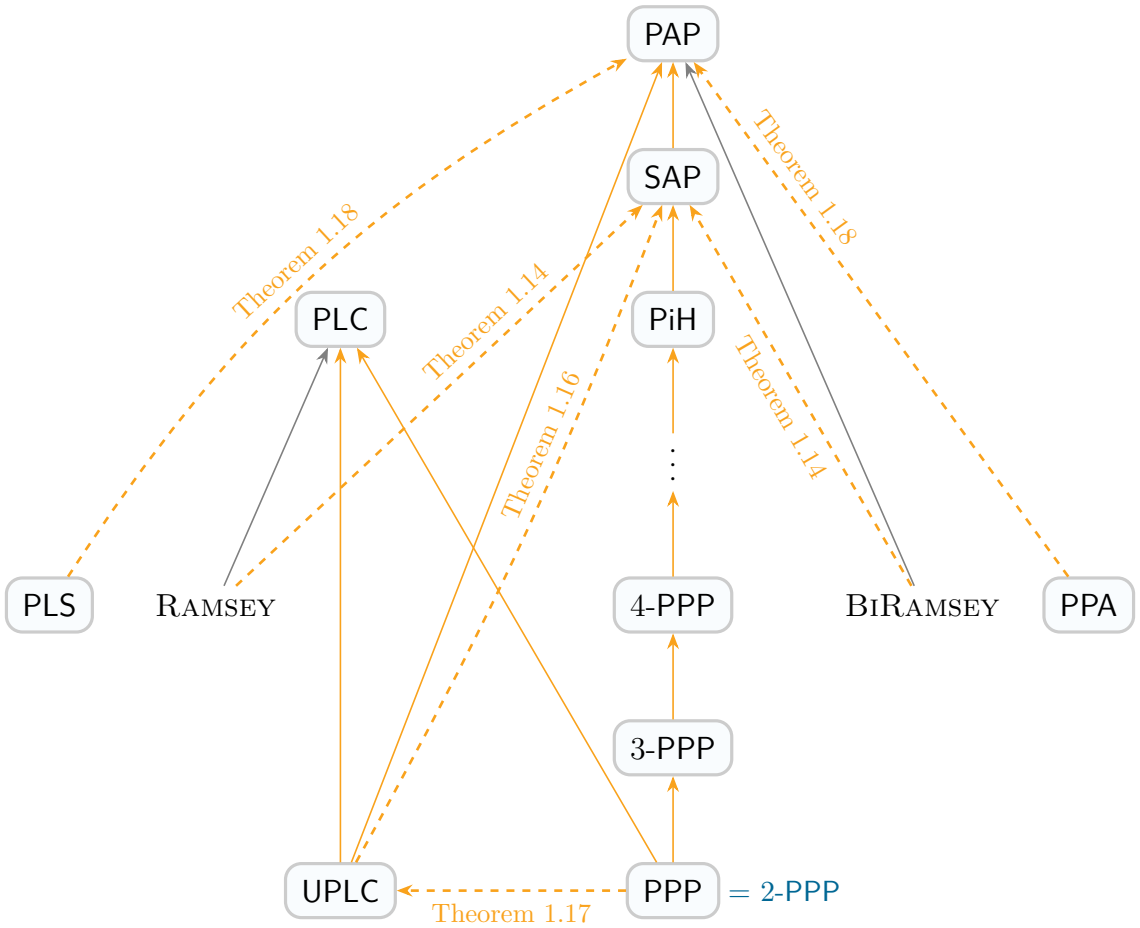

Figure 1: Complexity classes and problems defined by Generalized Pigeonhole Principles and Ramsey. An arrow $\mathsf{A} \to \mathsf{B}$ means $\mathsf{A} \subseteq \mathsf{B}$. An orange arrow $\mathsf{A} \to \mathsf{B}$ means $\mathsf{A} \subseteq \mathsf{B}$ and $\mathsf{B} \not\subseteq \mathsf{A}$ in the black-box setting. A dashed arrow $\mathsf{A} \dashrightarrow \mathsf{B}$ means $\mathsf{A} \not\subseteq \mathsf{B}$ in the black-box setting. All black-box separations in this figure are contributions of this work, labelled by the corresponding theorem.
We refer to the tower of t-PPP classes as the Pecking Order, while the union of t-PPP for all constant t is referred to as the Pigeon Hierarchy (PiH). We note that while our result BiRamsey $\not\subseteq$ SAP (in the black-box setting) applies for the standard parameter regime, BiRamsey $\in$ PAP is for a slightly smaller parameter.

search heuristics, is *also* not contained in PPP in the black-box setting. Using our techniques, we improve these results, showing that neither of these classes are even contained in PAP — the highest level of the Pecking Order:

Theorem 1.18. $\mathsf{PLS}^{dt} \not\subseteq \mathsf{PAP}^{dt}$ and $\mathsf{PPA}^{dt} \not\subseteq \mathsf{PAP}^{dt}$.

To prove these results we crucially employ our proof complexity techniques, combined with a generalization of the technical notion of *gluability* [GHJ+22], used in the previous black-box separations between PPP and PLS. We refer to Section 6 for technical details.

1.1.5 Quantum Complexity and the Pecking Order

Finally, we make an observation relating the Pecking Order to a problem in quantum complexity. It has remained an open problem to get a separation between BQP and BPP with respect to a *random oracle*, explicitly posed by Aaronson and Ambainis [AA14]. They even identified a technical barrier to proving such a separation, called the Aaronson-Ambainis conjecture, which has evaded attacks by experts in Boolean function analysis for almost a decade.

Recently, a breakthrough result was shown in this area by Yamakawa-Zhandry [YZ22] who obtained such a separation for *search* problems, evading the technical barrier that existed for decision problems. They also noted that their problem can be modified to show a separation between quantum and randomized query complexity in TFNP. We prove the totality of their problem using the averaging principle, and thus place it in PAP. This gives the first inclusion of the Yamakawa-Zhandry's Problem in a natural subclass of TFNP.

Theorem 1.19. Yamakawa-Zhandry's Problem is contained in PAP.

Furthermore, the proof of Theorem 1.19 suggests that the Yamakawa-Zhandry's problem necessarily corresponds to finding a poly(n)-collision. Given our structure theorem of the Pecking Order (Theorem 1.11), one may suspect that Yamakawa-Zhandry's problem is not contained in any lower level of the Pecking Order, e.g., PPP.

Finally, we remark that PAP is a rather loose upper bound for the Yamakawa-Zhandry's problem. The *structure* of the Yamakawa-Zhandry's problem, imposed by a specific choice of *error-correcting code* (ECC) in its definition, is lost during our reduction. The structure of the ECC is essential for its quantum speed-up, considering the fact that even PWPP cannot be solved efficiently by the quantum algorithm in the black-box setting [AS04].

1.2 Our Techniques

Our main separation results all follow from our black-box separations between $(t+1)$-$\textsc{Pigeon}_N^M$ and t-$\textsc{Pigeon}_{N'}^{M'}$, which in turn are proved by developing lower bound tools in *propositional proof complexity*. The connection between logic and TFNP has long been acknowledged [BK94, MPW00, Mor01, Tha02, BM04, Jeř09, KoNT11, BJ12, ST11, BB14, Jeř16, KoT22], and the first paper to prove oracle separations for TFNP [BCE+98] also invoked a Nullstellensatz lower bound for their result. Recently, *equivalences* have been established between complexity measures of certain proof systems and the complexity of black-box reductions to corresponding subclasses of TFNP. The first example of this was Göös, Kamath, Robere and Sokolov [GKRS19], followed by Göös et al. [GHJ+22] who proved equivalences for the prior natural-studied subclasses of TFNP and used these characterizations to provide the final remaining black-box separations between these subclasses. This was followed by Buss, Fleming and Impagliazzo [BFI23] who showed that this equivalence

holds for every TFNP problem — in a certain sense — although the corresponding proof system they construct is somewhat artificial, and hard to analyze directly.

Our techniques follow in this vein, as our lower-bound tools are directly inspired from propositional proof complexity. However, a major place where we depart is that there is currently no known natural proof system characterizing *any* of the classes in the Pecking Order, including PPP! This means that we do not have any lower-bound tools from proof complexity to borrow directly, and must instead develop new ones.

The tools we develop to prove our lower bounds are generalizations of *pseudoexpectation operators*, which have been instrumental in proving lower bounds for both the Sherali-Adams and the Sum-of-Squares proof system [FKP19]. Roughly speaking, a pseudoexpectation operator is an operator on polynomials that is indistinguishable from a true expectation operator, provided that we are only allowed to examine *bounded moments* of the distribution. In particular, it is possible to define pseudoexpectation operators that can *fool* a bounded adversary into thinking that there is a distribution over inputs to a total search problem R that do not witness any solutions (which is, of course, absurd). These operators are instrumental for understanding both Sherali-Adams and the Sum-of-Squares proof system, and have broad connections to the theory of approximation algorithms (see [FKP19] and the references therein for further details). Thus, our paper provides a new link between pseudoexpectation operators and the complexity of the *pigeonhole principle*.

Concurrent work of Fleming, Grosser, Pitassi, and Robere [FGPR24] introduced a new type of pseudoexpectation operator called a *collision-free pseudoexpectation* that is tailored for proving lower bounds against black-box PPP. Collision-freedom is a technical condition that is difficult to summarize without introducing a significant amount of notation (cf. Section 4). For now, we will say that the original notion of collision-free pseudoexpectation operators only applied to reductions to 2-PIGEON_N^{N+1} — that is, only for collisions of two pigeons, in the tight compression regime from $N+1$ pigeons to N holes. In our work, we introduce a broad generalization of collision-freedom that works for both an *arbitrary* number of collisions *and* an arbitrary compression rate. All of our main lower bound results follow designing generalized collision-free pseudoexpectations, combined with reductions placing various other problems inside the Pecking Order. To be specific, for separations *between* classes inside the Pecking Order (cf. Theorem 1.11), we design collision-free pseudoexpectations directly and prove that they satisfy the required properties. For our separations between the classes PLS and PPA, we provide a more general approach. Instead, we consider a general property (*gluability*, introduced by [GHJ+22]), and show that any pseudoexpectation for a problem satisfying this property is *automatically* collision-free for very strong parameters. This is a strong technical improvement (and conceptually different) application of gluability than had been applied by the earlier work of [GHJ+22]. We refer to Section 4 and Section 6 for further details on our pseudoexpectation technique.

1.3 Related Works

In this section we discuss the relationship of our work with concurrent work. We also discuss the how our black-box separations fit in with related work in cryptography.

Relationship with the work of Fleming, Grosser, Pitassi, and Robere [FGPR24]. In Section 4 we introduce a new type of pseudoexpectation operator called a (d, t, ε)-*collision-free pseudoexpectation* (cf. Definition 4.6), and show that the construction of such operators implies lower bounds for t-PPP^{dt}. This definition is a generalization of the notion of collision-free pseudoexpectation operators introduced in the concurrent work of Fleming, Grosser, Pitassi, and Robere. (To be precise, the original notion of a collision-free pseudoexpectation corresponds to the parameter setting $(d, 2, 1)$ in

the above definition). The lower bounds presented in this paper are orthogonal to the results of [FGPR24]. The main result in the concurrent work is to give a black-box separation between PPP^{dt} and its Turing closure $\mathsf{FP}^{\mathsf{PPP}^{dt}}$, a result which is incomparable to the results proved in our work.

Relationship with the work of Li [Li24]. This work precedes the work of Li [Li24], which generalizes the proof of Theorem 1.13 into a framework that separates two types of TFNP problems. In particular, [Li24] defines a new (semantic) subclass of TFNP called TFAP to captures TFNP problems with *abundant* solutions. [Li24] also introduces a notion called *semi-gluablility*. The main result in [Li24] is that there is no black-box reduction from any "semi-gluable" problems to any TFAP problems, and such separation could be extended to randomized reductions in most cases.

In this paper, we present a direct proof of Theorem 1.13 in Section 5, which is simpler than the proof of the main statement in Li [Li24], while sharing several key ideas with Li [Li24]. Related work in bounded arithmetic has been done by Müller [Mül21].

Relationship with the work of Bennett, Ghentiyala, and Stephens-Davidowitz [BGS24]. The concurrent work [BGS24] defines the classes in the Pecking Order using different notation, and prove results including black-box separations. We view our black-box separations (Theorem 1.11) as conceptually stronger, since the statement does not depend on the compression rate of the two problems. In contrast, the black-box separation from Bennett et al. builds upon the work of Rothblum & Vasudevan [RV22] and their technique only works with certain compression rates. For example, their technique cannot be used to infer lower bounds for Ramsey. They also study the relationship between the Pecking Order and certain coding and lattice problems, which has no overlap with our work.

The cryptographic picture. The search problems corresponding to the weak pigeonhole principles naturally show up in cryptography: given a hash function which compresses the input, how hard is it to find a collision? This is known as collision resistance hash function (CRH), which essentially corresponds to the average-case hardness of class PWPP. Recently there has been work investigating a generalization of this to *multi-collision resistance hash function (MCRH)* [Jou04, BKP18, KNY18, RV22], which naturally corresponds to n-PWPP defined in the Pecking Order. We refer readers to [BGS24] for a more comprehensive list of references.

It has remained open to show that there is no *fully black-box construction*[4] of CRH using MCRH.[5]. Our separation results in the Pecking Order (Theorem 1.11) make the first step towards showing a *fully black-box separation* between MCRH and CRH.

Our randomized separation of PPP from n-PWPP (Theorem 1.13) may also be of interest from a cryptography perspective. Berman et al. [BDRV18] and Komargodski et al. [KNY18] showed that there is no fully black-box construction of MCRH (corresponds to t-PWPP) using one-way permutation (corresponds to PPP). Technically, their separation result is not comparable to ours.

Both [BDRV18] and [KNY18] use an indirect approach to rule out the fully black-box construction. In particular, they present a fully black-box construction of *constant-round* statistically hiding commitment schemes using MCRH, and combine with a lemma from Haitner et al. [HHRS15] showing that there is no fully black-box construction of constant-round statistically hiding commitment

[4]Roughly speaking, a *fully black-box construction* of primitive A using primitive B in the cryptography setting is a uniform randomized black-box Turing reduction from the computational problem of B to the computational problem of A in average-case.

[5]To the best of our knowledge, this is still an open problem. We were notified by the authors of [KNY18] that there was an error in their proof that claimed to give a fully black-box separation between MCRH and CRH.

schemes from one-way permutation. In contrast, our proof is a direct one, which may offer more insights into the distinction between the one-way permutation and the MCRH.

1.4 Open problems

In our opinion, a fine-grained study of extremal combinatorics problems with respect to the two generalizations of PPP (PLC and Pecking Order) merits further research. The connections to cryptography, and quantum complexity also raise several intriguing questions. Here we list some open problems.

1. Can we strengthen our black-box separations to further rule out Turing reductions? Metamathematically, this could be interpreted as the (standard) pigeonhole principle is not strong enough for proving the Ramsey theorem.[6]
2. Can we show a black-box separations between PAP and PLC? This would separate the power of the "iterated pigeonhole" argument captured by PLC from the generalized pigeonhole principle.
3. Does RAMSEY lie in PAP? Recall that with a slight loss in parameter BIRAMSEY lies in PAP (Theorem 1.15).
4. The definition of PWPP is robust to different compression rate [Jeř16], e.g., both 2-PIGEON_N^{2N} and 2-$\text{PIGEON}_N^{N^2}$ are PWPP-complete. However, this is not known for t-PWPP for any $t > 2$ [Jou04]. Partial progress has been made in [BKP18, BGS24]. Either prove that these classes collapse under different compression rates or prove black-box separations.
5. Is UPLC $= n$-PWPP? Given Lemmas 6.15 and 6.16, this would be true if n-PWPP is robust for a certain range of compression rates.
6. Can we extend our separation results in Pecking Order to a fully black-box separation between the multi-collision-resistant hash function (MCRH) and the standard collision-resistant hash function (CRH)?
7. Is there a natural TFNP subclass — simpler than the Yamakawa-Zhandry's problem — in the Pecking Order which is contained in Total Function BQP?

Paper Organization In Section 2 we introduce necessary preliminaries for the paper, including definitions of the main TFNP search problems under consideration, the necessarily definitions for black-box TFNP, and the proofs of two basic structural properties of the Pecking Order. In Section 3, we characterize the complexity of RAMSEY and BIRAMSEY by proving Theorem 1.1 and relating them to the Pecking Order. Then, in Section 4, we introduce and develop our lower-bound technique of (d, t, ε)-collision-free pseudoexpectations, and use it to prove Theorem 1.11, the main structure theorem of the Pecking Order. In Section 5, we prove the randomized separations between PPP and any problem in Pecking Order that has non-trivial compression rate. Finally, Section 6 proves our other inclusions in and separations from the Pecking Order, including proving that $\text{PLS}^{dt}, \text{PPA}^{dt} \not\subseteq \text{PAP}^{dt}$, our results relating to PLC and UPLC, and also our results about the Yamakawa-Zhandry's problem.

[6] For the interested reader, we can state this formally in the language of bounded arithmetic as follows: RAMSEY is not uniformly Turing-reducible to PIGEON in the black-box model if and only if Ramsey's Theorem is logically independent from the Pigeonhole Principle over the bounded arithmetic theory $\forall \mathsf{S}_2(\mathsf{PV}(\alpha))$ [Mül21].

2 Preliminaries

Unless stated otherwise, $N = 2^n$. We also assume all the integer-valued functions defined in this paper are monotone and can be calculated efficiently.

2.1 Ramsey Theory

Ramsey theory is most generally defined to be a branch of combinatorics that studies how large some structure must be such that a property holds, which is usually the presence of a substructure. This branch was started by the study of the *Ramsey number*, which we define below.

Definition 2.1. $R(s,t)$ is defined to be the smallest integer n such that for any graph on n vertices, there is an independent set of size s or a clique of size t.

Theorem 2.2 (Ramsey's theorem [Ram29])**.** $R(s,t)$ is finite for every pair of integers s,t.

Finding upper and lower bounds on $R(s,t)$ has been a big program in the combinatorics community. But recently there has been interest in building *explicit* graphs with no clique or independent set of size K (which witness lower bounds for $R(K,K)$). These are motivated by connections to pseudorandomness. Constructing a pseudorandom object called a two-source *disperser* is essentially equivalent to constructing explicit bipartite Ramsey graphs, and getting better parameters has seen a long line of work [CG88, Coh16, CZ19, Li23]. This program makes progress on answering an old question of Erdős [Erd47]: can we construct explicit $O(\log n)$-Ramsey graphs on n vertices? We formally define this object below.

Definition 2.3 ($\boldsymbol{K}$-Ramsey)**.** A graph on N vertices is said to be K-Ramsey if it does not contain a clique or independent set of size K.

Definition 2.4 (Bipartite $\boldsymbol{K}$-Ramsey)**.** A bipartite graph on (N,N) vertices is said to be K-BiRamsey if it does not contain a biclique or independent set of size (K,K).

We recall a recent result of Li [Li23].

Theorem 2.5 ([Li23])**.** There exists a constant $c > 1$ such that for every integer N there exists a (strongly) explicit construction of a bipartite K-Ramsey graph on N vertices with $K = \log^c N$.

We now define TFNP problems corresponding to these extremal combinatorics results.

Definition 2.6 ($\boldsymbol{K}$-Ramsey)**.** For this problem, the input specifies a graph on $N = 2^n$ vertices; the parameter $K(n)$ is a function satisfying $R(K,K) \leq N$.

Input A pair (n, C), where $C : \{0,1\}^n \times \{0,1\}^n \mapsto \{0,1\}$ is a circuit of poly(n) size.

Solutions We have two kinds of solutions. One is a certificate that C is not a valid **encoding** of a simple undirected graph: u for which $C(u,u) = 1$ (self-loops) or u, v for which $C(u,v) \neq C(v,u)$ (directed edge). Else, we want to find K indices $V = \{v_1, v_2 \dots v_K\}$ which form a **clique** or **independent set**.

Note that we must choose K such that $R(K,K) \leq N$ in order for the problem to be total. An analogous condition holds for the bipartite analogue defined below below. We use $R_b(s,t)$ to denote the bipartite analogue of the Ramsey number.

Definition 2.7 ($\boldsymbol{K}$-BiRamsey)**.** For this problem, the input specifies a bipartite graph with N vertices on each side, where $N = 2^n$; the parameter $K(n)$ is a function satisfying $R_b(K,K) \leq N$.

Input A pair (n, C), where $C : \{0,1\}^n \times \{0,1\}^n \mapsto \{0,1\}$ is a circuit of poly(n) size.

Solutions (K, K) indices $U = \{u_1, u_2 \dots u_K\}$ and $V = \{v_1, v_2 \dots v_K\}$ such that (U, V) forms either a **biclique** or an **independent set**.

We drop the parameter K when we mean $K = n/2$. This is the existence result (up to logarithmic factors) guaranteed by the original theorem by Ramsey [Ram29]. Note that recently the first constant factor improvement in this existential result was obtained by Campos et al. [CGMS23], which corresponds to an exponential improvement in the Ramsey number. We also omit the size parameter N when it is clear in the context.

2.2 Decision Tree TFNP

Definition 2.8. A *query total search problem* is a sequence of relations $\mathrm{R} = \{R_n \subseteq \{0,1\}^n \times O_n\}$, where O_n are finite sets, such that for all $x \in \{0,1\}^n$ there is an $o \in O_n$ such that $(x, o) \in R_n$. A total search problem R is in TFNP^{dt} if for each $o \in O_n$ there is a decision tree T_o with depth poly$(\log n)$ such that for every $x \in \{0,1\}^n$, $T_o(x) = 1$ iff $(x, o) \in \mathrm{R}$.

In general, we use $DT(\mathrm{R})$ to denote the query complexity of a search problem R. To simplify the presentation and the relationship between the black-box model and the white-box model, we adhere to several conventions:

- For problems with a non-binary input alphabet, we simulate it with the usual binary encoding. For instance, in the pigeonhole principle, a mapping of a pigeon to a hole (i.e. a pointer in the range $[n]$) can be simulated by a $\log n$ bit binary encoding. All problems discussed in this paper have poly(n) alphabet size, and can therefore be simulated with $O(\log n)$ overheads.
- The problem R_n is permitted to have input bits on the order of poly(n).
- We sometimes abuse the notation by calling an individual relation R_n a search problem, rather than a whole sequence $\mathrm{R} = (R_n)$.
- The superscript "dt" is omitted when referring to TFNP^{dt} problems if the context makes it clear.

For instance, when t-$\textsc{Pigeon}_N^M$ problem is defined in the black-box model, the polynomial-size circuit h encoding the mapping of pigeons to holes is replaced by an oracle that queries the pigeons and maps them to holes. In the black-box model we also have an appropriate version of reducibility between search problems, where the reductions are computed by low-depth decision trees. We introduce this next.

Definition 2.9 (Decision Tree Reduction)**.** A *decision tree reduction* from relation $R \subseteq \{0,1\}^n \times O$ to $Q \subseteq \{0,1\}^m \times O'$ is a set of depth-d decision trees $f_i : \{0,1\}^n \to \{0,1\}$ for each $i \in [m]$ and $g_o : \{0,1\}^n \to O$ for each $o \in O'$ such that for any $x \in \{0,1\}^n$,

$$(x, g_o(x)) \in R \Leftarrow (f(x), o) \in Q,$$

where $f(x) \in \{0,1\}^m$ has $f_i(x)$ as the i-th bit. The *depth* of the reduction is d, and the *size* of the reduction is $\log m$. The *complexity* of the reduction is $d + \log m$, and we write $Q^{dt}(R)$ to denote the minimum complexity of a decision-tree reduction from R to Q, or ∞ if one does not exist.

We extend these notations to sequences in the natural way. If R is a single search problem and $\mathrm{Q} = (Q_m)$ is a sequence of search problems, then we denote by $\mathrm{Q}^{dt}(R)$ the minimum of $Q_m^{dt}(R)$ over all m. If $\mathrm{R} = (R_n)$ is also a sequence, then we denote by $\mathrm{Q}^{dt}(\mathrm{R})$ the function $n \mapsto \mathrm{Q}^{dt}(R_n)$. A TFNP^{dt} problem R can be *black-box reduced* to TFNP^{dt} problem Q, written $\mathrm{R} \leq_m \mathrm{Q}$, if there is a poly$(\log(n))$-complexity decision-tree reduction from R to Q.

In general, if a syntactic TFNP subclass A is defined by polynomial-time black-box reductions to a complete problem A, we will use A^{dt} to denote the class of query total search problems that can be efficiently black-box reduced to A. The next theorem motivates the decision-tree setting: constructing separations in the decision tree setting implies the existence of generic oracles separating the standard complexity classes.

Theorem 2.10 ([BCE+98], Informal)**.** For two syntactical TFNP subclasses A, B, $\mathsf{A}^{dt} \not\subseteq \mathsf{B}^{dt}$ *implies* the existence of a (generic) oracle O separating A from B, i.e., there is no black-box reduction from A to B.

We will also deal with *randomized* reductions in TFNP. We define these formally below.

Definition 2.11 (Randomized Decision Tree Reduction)**.** A *randomized decision tree reduction* from relation $R \subseteq \{0,1\}^n \times O$ to $Q \subseteq \{0,1\}^m \times O'$ is a distribution D of depth-d decision tree reductions from R to Q, such that for any $x \in \{0,1\}^n$,

$$\Pr_{((f_i),(g_o))\sim D}[(x, g_o(x)) \in R \Leftarrow (f(x), o) \in Q] \geq \frac{1}{2}.$$

2.3 Basic Structural Properties of the Pecking Order

By a simple padding argument, t-PPP forms a hierarchy when t increases.

Lemma 2.12. t-PPP $\subseteq (t+1)$-PPP. More generally, a-PPP $\subseteq b$-PPP if $a(n) \leq b(n)$ for all n.

Proof. Given a t-Pigeon instance (n, h) ($(t-1)N+1$ pigeons and N holes), we reduce it to a $(t+1)$-Pigeon instance (n, h') by adding N dummy pigeons. The i-th dummy pigeon is mapped to the i-th hole for any $i \in [N]$; the mapping for other pigeons are left unchanged. It is easy to verify that any $(t+1)$-collision in h' is a t-collision of h.

The general case follows from the same trick of adding dummy pigeons. □

The same proof can be generalized to show that a-PWPP $\subseteq b$-PWPP when $a(n) \leq b(n)$.

Definition 1.9. A function $b(n)$ is *polynomially close* to $a(n)$ if there exists a polynomial $p(n)$ such that for any n, $b(n) \leq a(p(n))$, and $a(n) \leq b(p(n))$.

Theorem 1.10. If a function $b(n)$ is polynomially close to $a(n)$, then a-PPP $= b$-PPP.

Proof. Let $p(n)$ be a polynomial such that $b(n) \leq a(p(n))$ for all n. Let $c(n) \coloneqq a(p(n))$. Lemma 2.12 implies that b-PPP $\subseteq c$-PPP. By the symmetry of the statement, it suffices to show that c-PPP $\subseteq a$-PPP.

Take a c-Pigeon instance $x \coloneqq (n, h)$. For technical convenience, we use an alternate encoding where $M = N \cdot (c(n) - 1)$, and the goal is to find either $(c(n)-1)$ pigeons mapped to 1 or a $c(n)$-collision. It is easy to show the equivalence between this alternate version and the original version.

We now construct an a-Pigeon instance $x' \coloneqq (h', n')$, where $n' = p(n)$. Let $N' \coloneqq 2^{n'}, M' \coloneqq a(n') \cdot N'$ be the number of holes and pigeons that are supposed to be in x'. x' is simply chosen as N'/N disjoint copies of x in parallel. It is easy to verify that the number of holes in x' is $N \cdot (N'/N) = N'$, and the number of pigeons in x' is

$$M \cdot (N'/N) = (c(n) - 1) \cdot N' = (a(n') - 1) \cdot N' = M'.$$

Formally, we index each hole in x' by a pair $(i,l) \in [N] \times [N'/N]$; each pigeon in x' is also indexed by a pair $(j,l) \in [M] \times [N'/N]$. Now, $h' : [M] \times [N'/N] \mapsto [N] \times [N'/N]$ is defined as

$$h'(j,l) := (h(j),l), \forall (j,l) \in [M] \times [N'/N].$$

Note that any collision in x' must come from a single copy of x, because different copies of x are disjoint. Also given that $a(n') = c(n)$, any solution of x' immediately corresponds to a solution of x. This concludes the correctness of our reduction.

□

One can easily generalize the proof of Theorem 1.10 to show that a-PWPP $=$ b-PWPP when $a(n), b(n)$ are polynomially close. Note that all polynomial functions are polynomially close to each other. Therefore, t-PPP is equivalent to PAP for any choice of $t(n) = \Theta(\text{poly}(n))$.

3 Ramsey and the Pecking Order

In this section, we show that problems in the Pecking Order can be reduced to Ramsey and BiRamsey (Lemma 3.1). As a consequence of the generality of parameters in the structural theorem of the Pecking Order (Theorem 1.11), we immediately get that Ramsey^{dt} is not in PPP^{dt} (Theorem 1.3). In the reverse direction, we show that BiRamsey with a slightly weaker parameter is included in PAP (Theorem 1.15); combined with Theorem 1.3, we get an almost tight characterization of BiRamsey using the Pecking Order.

We start with reducing problems in the Pecking Order to Ramsey, which directly generalize a reduction that appeared in [KNY19]. For ease of discussion, we only state the following result in terms of Ramsey, yet the same argument also holds true for BiRamsey.

Lemma 3.1 (Generalization of [KNY19, Theorem 3.1])**.** Suppose there exists a graph on N vertices with no K clique or independent set. Then t-Pigeon$_N^M$ can be black-box reduced to $t(K-1)$-Ramsey when $2t(K-1) \leq \log M$.

Proof. Given a graph $G_0(V_0, E_0)$ on N vertices such that it does not contain a K clique or independent set, we will build an instance G given by $(\log M, E)$ of $t(K-1)$-Ramsey from an instance C of t-Pigeon$_N^M$.

Let $h : [M] \mapsto [N]$ be an instance of t-Pigeon$_N^M$. We will define a $t(K-1)$-Ramsey instance G using the *graph hash product* from [KNY19]. Let $G = G_0 \otimes h = (V, E)$ be a graph on $V = [M]$ such that

$$(u,v) \in E \iff h(u) = h(v) \text{ or } (h(u), h(v)) \in E_0.$$

Now we prove that the solution S returned by $t(K-1)$-Ramsey witnesses a t-collision in h. Let S' be set $\{h(u) \mid u \in S\} \subseteq V(G_0)$. By the definition of G, S' is a clique if S is a clique, and S' is an independent set if S is an independent set. Therefore, S' is either a clique or an independent set.

Since G_0 does not contain a K clique or independent set, we have $|S'| < K$. Given $|S| = t(K-1)$, by an averaging argument S must witness a $\frac{t(K-1)}{K-1} = t$ collision in h.

□

Using the probabilistic method, Erdös [Erd47] shows that there exists a graph on N vertices with no clique or independent set of size $K = 2\log N$. This gives us the following theorem as a corollary:

Theorem 1.3. Whenever $M \geq N^{4t}/4^t$, t-Pigeon$_N^M$ can be black-box reduced to Ramsey and its bipartite variant BiRamsey.

We also note that by Theorem 2.5, we have that for some $c > 1$, we have an explicit efficient black-box reduction from t-PIGEON$_N^M$ to RAMSEY whenever $2t(\log^c N - 1) \leq \log M$. As a consequence of these reductions, we get two corollaries of Theorem 1.11 regarding the place of K-RAMSEY in TFNP. Below, we assume Theorem 1.11, which is proved in Section 4.

Theorem 1.14. RAMSEY^{dt}, BIRAMSEY$^{dt} \notin$ SAP^{dt}. In particular, RAMSEY^{dt}, BIRAMSEY$^{dt} \notin$ PPP^{dt}.

Proof. Theorem 1.3 shows that there exists some polynomial $p(n)$ such that $p(n)$-PIGEON$_N^M$ reduces to RAMSEY$_M$ for large enough M. However, by Theorem 1.11, we know that there's no black-box reduction from $p(n)$-PIGEON$_N^M$ to $t(n)$-PIGEON$_{N'}^{M'}$ for any subpolynomial $t(n)$. This implies that RAMSEY is not in $t(n)$-PPP^{dt} for any subpolynomial $t(n)$. □

Actually, by being a little more careful with the parameters, we can prove a stronger version of the theorem following the same argument, which shows that even $\log^c(N)$-RAMSEY^{dt} is not in SAP^{dt}.

Finally, we show that with a slight loss in the parameter, BIRAMSEY fits into the Pecking Order.

Theorem 1.15. (Generalization of [KNY19, Theorem 3.10]) $\frac{n-\log n}{2}$-BIRAMSEY $\in$ PAP.

Proof. Given a bipartite graph $G = ([N], [N], E)$, we say $E(x, y) = 1$ if $(x, y) \in E$ and $E(x, y) = 0$ otherwise. We construct an n-PIGEON$_{N/n}^N$ instance using function $h : [N] \mapsto \{0, 1\}^{n-\log n}$ defined by

$$h(y) := (E(x, y))_{x \in [n-\log n]}.$$

Our goal is to prove that we can efficiently find a clique or independent set of size $(n - \log n)/2$ from an n-collision in h. Let $y_1, \ldots, y_n \in [N]$ be an n-collision in h. Then by the definition of h, we have for each $x \in [n - \log n]$, $E(x, y_1) = E(x, y_2) = \cdots = E(x, y_n)$.

At least half of the x's in $[n - \log n]$ will have the same value for $E(x, y_1)$, and we let these indices be $x_1 < \ldots < x_{(n-\log n)/2}$. This gives us that for all $i, j \in [(n - \log n)/2]$, $E(x_i, y_j) = E(x_1, y_j) = E(x_1, y_1)$. Therefore, $(\{x_1, \ldots, x_{(n-\log n)/2}\}, \{y_1, \ldots, y_{(n-\log n)/2}\})$ is either a biclique or an independent set.

□

Query complexity of Ramsey. Besides the relative complexity of RAMSEY, one might also ask questions about its query complexity in various models (deterministic, randomized, quantum). To find a clique or independent set of size k, it certainly suffices to query all the vertices of an arbitrary subgraph of size $R(k, k)$. This gives us an upper bound on the deterministic query complexity of RAMSEY of $\binom{R(n/2,n/2)}{2}$. Plugging in the best known upper bound on the diagonal Ramsey number [CGMS23], we get an upper bound on $N^{2-\epsilon}$ for a small constant ϵ. On the lower bound front, we can infer that the quantum query complexity of RAMSEY is at least $N^{1-o(1)}$ due to the reduction in Lemma 3.1 combined with the quantum query lower bound by Liu and Zhandry [LZ19]. Since the best lower bound we have on $R(t, t)$ is $2^{\widetilde{O}(t/2)}$ [Erd47], improving even the deterministic lower bound significantly beyond $\Omega(N)$ would give a *new combinatorics result*! The best deterministic lower bound known is due to Conlon et al. [CFGH19]. We note this as an exciting approach to getting better lower bounds for the Ramsey numbers.

4 Structure of the Pecking Order

In this section we prove Theorem 1.8, our separation of the Pigeon Hierarchy, restated here for convenience.

Theorem 1.8. If t is a constant and M, N, M', N' are parameters chosen so that $M' = \mathrm{poly}(M)$, $M \geq tN + 1$, $M' \geq (t-1)N' + 1$, then $(t+1)$-$\textsc{Pigeon}_N^M$ does not have an efficient black-box reduction to t-$\textsc{Pigeon}_{N'}^{M'}$. In particular, $(t+1)$-$\mathsf{PPP}^{dt} \not\subseteq t$-$\mathsf{PPP}^{dt}$ for any constant t, and so PiH^{dt} forms a strict hierarchy in the black-box setting.

As we discussed in the introduction, our lower bounds are proved using (generalizations of) tools from propositional proof complexity, and in particular the theory of *pseudoexpectation operators*. Our main theorem is proved by generalizing the notion of *collision-free pseudoexpectation operators*, designed by Fleming, Grosser, Pitassi, and Robere to give a black-box separation between PPP^{dt} from its Turing-closure, to the entire Pecking Order [FGPR24].

Pseudoexpectation Operators. First we introduce the notion of a *pseudoexpectation operator*, and for this we need to recall some basic results about multilinear polynomials. Let $x_1, x_2, \ldots, x_n$ be a family of $\{0,1\}$-valued variables. We consider real-coefficient polynomials $p \in \mathbb{R}[x_1, \ldots, x_n]$ over these variables. All polynomials that we consider are *multilinear*, meaning the individual degree of any variable is at most 1. The algebra of multilinear polynomials is described by the quotient ring $\mathbb{R}[x_1, \ldots, x_n]/\langle x_i^2 - x_i \rangle_{i=1}^n$. Formally, the addition of two multilinear polynomials is still a multilinear polynomial, but, if we multiply two multilinear polynomials then, after multiplication, we drop the exponents of all variables to 1. For example, $(x+y)(x+y) = x + 2xy + y$, as multilinear polynomials.

For any $S, T \subseteq [n]$ with $S \cap T = \emptyset$, define the polynomial

$$C_{S,T} := \prod_{i \in S} x_i \prod_{j \in T} (1 - x_j).$$

Note that for $\{0,1\}$-assignments the polynomial $C_{S,T}$ encodes the truth value of the conjunction $\bigwedge_{i \in S} x_i \wedge \bigwedge_{j \in T} \overline{x}_j$, and thus we will refer to $C_{S,T}$ as a "conjunction" in an abuse of notation. If $R_n \subseteq \{0,1\}^n \times O$ is a query total search problem, and $o \in O$, then a conjunction C *witnesses* the solution o if $C(x) = 1 \Rightarrow (x, o) \in R_n$ for every $x \in \{0,1\}^n$. Similarly, we say a conjunction C *witnesses* R_n if it witnesses some solution to R_n. Each conjunction $C = C_{S,T}$ is naturally associated with a partial restriction $\rho(C) \in \{0, 1, *\}$ defined by setting

$$\rho(C)_i = \begin{cases} 1 & i \in S, \\ 0 & i \in T, \\ * & \text{otherwise.} \end{cases}$$

Definition 4.1. Let $n \geq d$ be positive integers. Let $\mathcal{P}_{n,d}$ be the collection of all degree $\leq d$ multilinear polynomials over variables $x_1, \ldots, x_n$. An operator $\widetilde{\mathbb{E}} : \mathcal{P}_{n,d} \to \mathbb{R}$ is a *degree-d pseudoexpectation operator* if it satisfies the following three properties:

- **Linearity.** $\widetilde{\mathbb{E}}$ is linear.
- **Normalized.** $\widetilde{\mathbb{E}}[1] = 1$.
- **Nonnegativity.** $\widetilde{\mathbb{E}}[C] \geq 0$ for all degree $\leq d$ conjunctions C.

Furthermore, if $R \subseteq \{0,1\}^n \times O$ is a query total search problem, then $\widetilde{\mathbb{E}}$ is *R-Nonwitnessing* if it additionally satisfies the following property:

- ***R*-Nonwitnessing.** $\widetilde{\mathbb{E}}[C] = 0$ for any conjunction C witnessing R.

If $\mathcal{F}$ is any sequence of degree $\leq d$ multilinear polynomials, then we write $\widetilde{\mathbb{E}}[\mathcal{F}] := \sum_{p \in \mathcal{F}} \widetilde{\mathbb{E}}[p]$.

Pseudoexpectation operators were originally introduced to prove lower bounds on the degree of Sherali-Adams refutations for unsatisfiable CNF formulas [FKP19]. Often the easiest way to construct a pseudoexpectation is to construct an object called a *pseudodistribution* instead. We introduce pseudodistributions next:

Definition 4.2. Let $x_1, \ldots, x_n$ be a set of $\{0,1\}$-valued variables, and let $d \leq n$. A *degree-d pseudodistribution* over these variables is a family of probability distributions

$$\mathcal{D} = \{\mathcal{D}_S : S \subseteq [n], |S| \leq d\},$$

such that the following properties hold:

- For each set $S \subseteq [n]$, $|S| \leq d$, $\mathcal{D}_S$ is supported on $\{0,1\}^S$, interpreted as boolean assignments to variables in $\{x_i : i \in S\}$.
- For each $S, T \subseteq [n]$, $|S|, |T| \leq d$, we have $\mathcal{D}_S^{S \cap T} = \mathcal{D}_T^{S \cap T} = \mathcal{D}_{S \cap T}$, where $\mathcal{D}_A^B$ for $B \subseteq A$ is the marginal distribution of variables in B with respect to $\mathcal{D}_A$.

The following standard lemma allows us to construct pseudoexpectations from pseudodistributions. In fact, the two objects are equivalent, but we will not need the converse construction in this paper.

Lemma 4.3 ([FKP19])**.** Let $\mathcal{D}$ be a degree-d pseudodistribution over variables $x_1, \ldots, x_n$. The operator $\widetilde{\mathbb{E}}$ defined by

$$\widetilde{\mathbb{E}}\left[\prod_{i \in S} x_i\right] = \Pr_{y \sim \mathcal{D}_S}[\forall i \in S : y_i = 1]$$

and extended to all multilinear polynomials by linearity, is a degree-d pseudoexpectation.

The following pseudodistribution, and its accompanying pseudoexpectation, is the central pillar of all lower bounds in this paper. This example is, in fact, one of the classical examples of a pseudodistribution [GM08, Lemma 2].

Definition 4.4 (Matching Pseudodistribution)**.** Consider t-$\textsc{Pigeon}_N^M$ with $M \geq (t-1)N + 1$ pigeons and N holes, and let $d \leq N/2$. The *degree-d matching pseudodistribution* for this instance is the following pseudodistribution. For any set of input variables S in t-$\textsc{Pigeon}_N^M$, let $p(S)$ denote the set of pigeons mentioned among variables in S. For each subset S of variables with $|S| \leq d$, define the distribution $\mathcal{D}_S$ by sampling a uniformly random matching from the pigeons in $p(S)$ to $|p(S)|$ holes, and assigning the variables in S according to this matching. Further define $\mathcal{D} = \{\mathcal{D}_S : |S| \leq d\}$ to be the collection of all such distributions.

The above construction indeed defines a pseudodistribution, as shown in [GM08]. To put it simply, if we sample a uniformly random matching from t pigeons to t holes and then marginalize one pigeon out, then the result is a uniformly random matching from $t-1$ pigeons to $t-1$ holes.

Lower Bounds for the Pecking Order. We are now ready to prove the main lower bound result of this paper. As we have mentioned above, degree-d pseudoexpectation operators were originally introduced to prove lower bounds for Sherali-Adams refutations. In order to do this, we must construct high-degree pseudoexpectation operators which are additionally *nonwitnessing*, meaning that $\widetilde{\mathbb{E}}[C] = 0$ whenever C is a conjunction witnessing a solution to our search problem R. A recent work [GHJ+22] has shown that a query total search problem R is in PPADS^{dt} if and only if an unsatisfiable CNF formula $\neg\text{Total}(R)$ expressing the totality of R has low-degree unary Sherali-Adams refutations.

This means that constructing a nonwitnessing pseudoexpectation for R automatically implies that $R \notin \mathsf{PPADS}^{dt}$. However, R can admit a nonwitnessing pseudoexpectation but still lie in PPP^{dt} or higher up in the Pecking Order — for example, the matching pseudoexpectation is a nonwitnessing pseudoexpectation for 2-PIGEON, which is the complete problem for PPP.

Therefore, to prove lower bounds for higher levels of the Pecking Order, we must strengthen the definition of a pseudoexpectation. To introduce this strengthening, we need the following auxiliary definition.

Definition 4.5. Let $R \subseteq \{0,1\}^n \times O$ be a query total search problem, and let $\mathcal{F}$ be a family of degree-d conjunctions over input variables of R. For $t \geq 2$, the family $\mathcal{F}$ is said to be *t-witnessing for R* if for any subset $\mathcal{S} \subseteq \mathcal{F}$ with $|\mathcal{S}| = t$, either $\prod_{C \in \mathcal{S}} C \equiv 0$, or $DT(R \restriction \rho) \leq d$, where ρ is the concatenation of $\rho(C)$ for all $C \in \mathcal{S}$, and D is some universal constant. In other words, either every subset of t conjunctions is inconsistent, or, there is a shallow decision tree solving the restricted problem $DT(R \restriction \rho)$, where ρ is the union of partial assignments corresponding to the t conjunctions. We say that $\mathcal{F}$ is *t-witnessing* if the problem R is clear from context.

Definition 4.6. Let $R \subseteq \{0,1\}^n \times O$ be a total query search problem, let d, t be positive integers, and let $\varepsilon > 0$ be a real parameter. Let $\widetilde{\mathbb{E}}$ be a degree $D \geq d$ pseudoexpectation operator. Then $\widetilde{\mathbb{E}}$ is *(d, t, ε)-collision-free for R* if it satisfies the following property:

- ***t*-Collision-Freedom.** $\widetilde{\mathbb{E}}[\mathcal{F}] \leq \varepsilon$, for every t-witnessing family $\mathcal{F}$ of degree $\leq d$ conjunctions.

The notion of a collision-free pseudoexpectation was introduced by Fleming, Grosser, Pitassi, and Robere [FGPR24] in the special case where $t = 2, \varepsilon = 1$, in order to separate PPP from its Turing closure in the black-box setting. The above definition generalizes this notion to arbitrary size-t collisions. As we will see, t-Collision Freedom is the additional property that is required of a pseudoexpectation in order to rule out membership of problems in t-PPP. The following theorem generalizes the same theorem for $M = N + 1, t = 2, \varepsilon = 1$, proved by [FGPR24].

Theorem 4.7. Let $R \subseteq \{0,1\}^n \times O$ be a query total search problem. Let M, N, t be positive integers with $M \geq (t-1)N + 1$, and let $0 \leq \varepsilon < M/N$ be any real parameter. If there is a (d, t, ε)-collision-free pseudoexpectation operator for R then there is no depth-d decision-tree reduction from R to t-PIGEON_N^M.

Proof. For the sake of intuition, we first prove this for the case of $(d, 2, \varepsilon)$-pseudoexpectation operators, corresponding to the classical TFNP class PPP, but the argument easily generalizes to arbitrary t, N, and $M \geq (t-1)N + 1$. Let us assume by way of contradiction that there is a degree-d decision-tree reduction from R to 2-PIGEON_N^{N+1} for some N, and let $\widetilde{\mathbb{E}}$ denote the $(d, 2, \varepsilon)$-collision-free pseudoexpectation for R. Let $T_1, T_2, \ldots, T_{N+1}$ denote the depth-d decision trees in the 2-PIGEON_N^{N+1} instance produced by the reduction mapping the pigeons to holes. First, for any decision tree T_i let $L(T_i)$ denote the leaves of T_i, and, for any leaf ℓ of T_i, let C_ℓ denote the conjunction obtained by multiplying the literals along the path to ℓ. An easy induction on the depth of the tree shows that for every tree T_i,

$$\sum_{\ell \in L(T_i)} \widetilde{\mathbb{E}}[C_\ell] = 1.$$

Now for any hole $h \in [N]$, let $\mathcal{F}_h$ denote the set of all conjunctions C_ℓ that correspond to paths of any decision tree among $T_1, \ldots, T_{N+1}$ such that the leaf ℓ of that path is labelled with h. Since this instance of 2-PIGEON_N^{N+1} is obtained via a depth-d reduction from R, for any pair of conjunctions $C, D \in \mathcal{F}_h$, either C and D are inconsistent, or, $DT(R \restriction \rho(CD)) \leq d$ since a collision of two pigeons

implies that we can recover a solution of R after at most d more queries. It follows that the family $\mathcal{F}_h$ is 2-witnessing, in the language of Definition 4.5.

Since $\mathcal{F}_h$ is 2-witnessing for each hole, it follows that $\widetilde{\mathbb{E}}[\mathcal{F}_h] \leq \varepsilon$ for every $h \in [N]$ since $\widetilde{\mathbb{E}}$ is $(d, 2, \varepsilon)$-collision-free. Now, observe that

$$\sum_{i=1}^{N+1} \sum_{\ell \in L(T_i)} \widetilde{\mathbb{E}}[C_\ell] = \sum_{h=1}^{N} \widetilde{\mathbb{E}}[\mathcal{F}_h],$$

since each of the $N+1$ pigeons are mapped to exactly one hole under a total assignment to the variables. This means that

$$N+1 = \sum_{i=1}^{N+1} \sum_{\ell \in L(T_i)} \widetilde{\mathbb{E}}[C_\ell] = \sum_{h=1}^{N} \widetilde{\mathbb{E}}[\mathcal{F}_h] \leq \varepsilon N < N+1,$$

which is a contradiction.

To generalize this for arbitrary t, we instead consider reductions to t-$\textsc{Pigeon}_N^M$ and substitute $(d, 2, \varepsilon)$-witnessing with (d, t, ε)-witnessing in the above proof. Since the instance of t-$\textsc{Pigeon}_N^M$ is obtained by depth-d reduction from R, it follows now that for every set of t distinct conjunctions $C_1, C_2, \ldots, C_t \in \mathcal{F}_h$, either $\prod_{i=1}^{t} C_i$ is inconsistent, or, $DT(R \restriction \rho(C_1 C_2 \cdots C_t)) \leq d$, since a collision of t pigeons implies that we can recover a solution of R after at most d more queries. Therefore, $\mathcal{F}_h$ is now t-witnessing, and so $\widetilde{\mathbb{E}}[\mathcal{F}_h] \leq \varepsilon$. We now have

$$M = \sum_{i=1}^{M} \sum_{\ell \in L(T_i)} \widetilde{\mathbb{E}}[C_\ell] = \sum_{i=1}^{N} \widetilde{\mathbb{E}}[\mathcal{F}_h] \leq \varepsilon N < M,$$

a contradiction. □

The previous theorem gives us a powerful method to prove lower bounds against levels of the Pecking Order. In particular, we will be able to show that the *matching pseudoexpectation* is t-collision-free against $(t+1)$-$\textsc{Pigeon}_N^M$ for *all* t. The main observation here is that the union of t matchings from M to N has a collision of size at most t, and therefore cannot witness a $(t+1)$-collision of pigeons. This means that any t-subset of conjunctions from a t-witnessing family of matchings is inconsistent. The next technical lemma shows how to upper-bound the weight of such inconsistent families.

Lemma 4.8. Let $x_1, \ldots, x_n$ be a set of variables, and let t, d be chosen so that $(t-1)d^2 \leq n$. Let $\mathcal{F}$ be any family of degree $\leq d$ conjunctions over these variables, such that for every subset $\mathcal{S} \subseteq \mathcal{F}$ with $|\mathcal{S}| = t$, $\prod_{C \in \mathcal{S}} C \equiv 0$. If $\widetilde{\mathbb{E}}$ is a degree $D \geq (t-1)d^2$ pseudoexpectation operator, then $\widetilde{\mathbb{E}}[\mathcal{F}] \leq t-1$.

Proof. Let us first observe that the statement of the lemma would obviously be true if $\widetilde{\mathbb{E}}$ were the expectation over a *true* probability distribution. This is because at most $t-1$ distinct conjunctions in $\mathcal{F}$ are consistent with any total assignment, and thus no set of t conjunctions can be simultaneously activated under a sample from the true probability distribution.

To prove this claim for $\widetilde{\mathbb{E}}$ we first need to introduce some notation. Let T be any decision tree querying the variables $x_1, \ldots, x_n$ and outputting 0 or 1. Let $L_b(T)$ denote the leaves of T labelled with $b \in \{0, 1\}$, and let $L(T) = L_0(T) \cup L_1(T)$. For any leaf $\ell \in L(T)$ let C_ℓ denote the conjunction of literals on the path from the root to ℓ, and let $\rho_\ell = \rho(C_\ell)$ denote the partial assignment corresponding to this path. If the depth of T is at most d, define

$$\widetilde{\mathbb{E}}[T] := \sum_{\ell \in L_1(T)} \widetilde{\mathbb{E}}[C_\ell].$$

An easy induction on the depth of T shows that $\widetilde{\mathbb{E}}[T] \leq 1$.

Starting with the family $\mathcal{F}$, we create a depth $\leq (t-1)d^2$ decision tree T such that

$$\widetilde{\mathbb{E}}[\mathcal{F}] \leq (t-1)\widetilde{\mathbb{E}}[T] \leq t-1.$$

If ρ is a partial assignment, then let $\mathcal{F} \restriction \rho = \{C \restriction \rho : C \in \mathcal{F}\}$, where it is understood that we remove any conjunctions that are set to 0 or 1 under the restriction ρ.

We construct the decision tree T by the following recursive algorithm. The decision tree maintains a partial assignment to the above variables. Initially, $\rho = \emptyset$. The algorithm proceeds in rounds. If $\mathcal{F} \restriction \rho = \emptyset$, we halt and output 1, if there are any conjunctions in $\mathcal{F}$ consistent with ρ, or 0 otherwise. We proceed assuming $\mathcal{F} \restriction \rho \neq \emptyset$. In this case, we choose $t-1$ conjunctions $C_1, C_2, \ldots, C_{t-1}$ in $\mathcal{F} \restriction \rho$ and query all unqueried variables among these conjunctions — if there are less than $t-1$, then we simply query all the variables among all remaining conjunctions. After this querying stage, we have learned a partial assignment σ to the newly queried variables, and we then recurse on the family $\mathcal{F} \restriction \rho\sigma$.

The construction above plainly terminates on every branch, since we are reducing the length of each conjunction after every round of querying. We argue that the depth of the tree is at most $(t-1)d^2$ and that $\widetilde{\mathbb{E}}[\mathcal{F}] \leq (t-1)\widetilde{\mathbb{E}}[T]$, which completes the proof of the theorem.

First, we argue that the depth of the tree is at most D. To see this, we observe that in each round we query at most $(t-1)d$ variables, and we argue that on every branch the algorithm terminates after at most d rounds. To see this, consider the i-th round, where we query conjunctions $C_1, \ldots, C_{t-1}$. Since $\mathcal{F}$ is t-witnessing, it follows that every conjunction C remaining in $\mathcal{F}$ must conflict with at least one literal contained in $C_1, \ldots, C_{t-1}$. Therefore, after querying all unqueried variables in $C_1, \ldots, C_{t-1}$, we must query at least one variable from every remaining conjunction in $\mathcal{F}$. This means that at the end of the ith round, we must reduce the length of every remaining conjunction by at least one. Since each conjunction has at most d variables to begin with, it follows that the entire process can proceed for at most d rounds. Thus, the depth of the tree is at most D.

Let us now see that $\widetilde{\mathbb{E}}[\mathcal{F}] \leq (t-1)\widetilde{\mathbb{E}}[T]$. First, we observe that since the depth of T is at most $(t-1)d^2$, it follows that every for every leaf ℓ of T the conjunction C_ℓ has degree at most D. This means that $\widetilde{\mathbb{E}}[T]$ is well-defined since $\widetilde{\mathbb{E}}$ is a degree-D pseudoexpectation.

So, it remains to show that $\widetilde{\mathbb{E}}[\mathcal{F}] \leq (t-1)\widetilde{\mathbb{E}}[T]$, noting that $\widetilde{\mathbb{E}}[T] \leq 1$ for any decision tree T. For any node u in T let T_u denote the subtree rooted at u, and let C_u denote the conjunction of the literals along the path to u. We prove by induction on the height of u that

$$\sum_{C \in \mathcal{F}} \widetilde{\mathbb{E}}[C_u C] \leq (t-1) \sum_{\ell \in L_1(T_u)} \widetilde{\mathbb{E}}[C_u C_\ell]. \tag{1}$$

Once we have this equation, setting u to be the root node r yields

$$\sum_{C \in \mathcal{F}} \widetilde{\mathbb{E}}[C] \leq (t-1) \sum_{\ell \in L_1(T)} \widetilde{\mathbb{E}}[C_\ell] = (t-1)\widetilde{\mathbb{E}}[T],$$

as desired.

If u is a 1-leaf node of T_u, then $L_1(T_u) = \{u\}$ and so

$$\sum_{\ell \in L_1(T_u)} \widetilde{\mathbb{E}}[C_u C_\ell] = \widetilde{\mathbb{E}}[C_u].$$

On the other hand, for any leaf u, by construction of the tree if $C_u C \not\equiv 0$ then $C_u C = C_u$. Since the family $\mathcal{F}$ is t-witnessing, there are at most $t-1$ possible choices of $C \in \mathcal{F}$ such that $C_u C \not\equiv 0$

since every set of t conjunctions in $\mathcal{F}$ are inconsistent. Therefore

$$\sum_{C\in\mathcal{F}} \widetilde{\mathbb{E}}[C_u C] \le (t-1)\widetilde{\mathbb{E}}[C_u] = (t-1)\sum_{\ell\in L_1(T_u)} \widetilde{\mathbb{E}}[C_u C_\ell],$$

proving the base case of (1).

For the inductive step, consider a node u in T querying a variable x_i with children v_0, v_1 corresponding to the two outcomes of the query. Then

$$\begin{aligned}
\sum_{C\in\mathcal{F}} \widetilde{\mathbb{E}}[C_u C] &= \sum_{C\in\mathcal{F}} \widetilde{\mathbb{E}}[C_u x_i C + C_u(1-x_i)C] \\
&= \sum_{C\in\mathcal{F}} \widetilde{\mathbb{E}}[C_u x_i C] + \widetilde{\mathbb{E}}[C_u(1-x_i)C] \\
&= \sum_{C\in\mathcal{F}} \widetilde{\mathbb{E}}[C_{v_1} C] + \sum_{C\in\mathcal{F}} \widetilde{\mathbb{E}}[C_{v_0} C] \\
&\le (t-1)\sum_{\ell\in L_1(T_{v_1})} \widetilde{\mathbb{E}}[C_{v_1} C_\ell] + (t-1)\sum_{\ell\in L_1(T_{v_0})} \widetilde{\mathbb{E}}[C_{v_0} C_\ell] \\
&= (t-1)\left(\sum_{\ell\in L_1(T_{v_1})} \widetilde{\mathbb{E}}[C_u x_i C_\ell] + \sum_{\ell\in L_1(T_{v_0})} \widetilde{\mathbb{E}}[C_u(1-x_i)C_\ell]\right) \\
&= (t-1)\sum_{\ell\in L_1(T_u)} \widetilde{\mathbb{E}}[C_u C_\ell],
\end{aligned}$$

where the inequality follows by the induction hypothesis, and the last equality follows since the leaves of T_u are exactly the union of leaves of T_{v_0} and T_{v_1}. This proves (1) for all nodes u of T, completing the proof. □

With this technical lemma in hand, we are now ready to prove the main theorem of this section.

Theorem 4.9. Let t, d, M, N, M', N' be positive integers chosen so that $M \ge tN+1$, $M' \ge (t-1)N'+1$, and $(t-1)d^2 \le N/2$. Then the $(t+1)$-$\textsc{Pigeon}_N^M$ problem does not have a depth-d decision-tree reduction to t-$\textsc{Pigeon}_{N'}^{M'}$.

Proof. Let $D = (t-1)d^2 \le N/2$, and let $\widetilde{\mathbb{E}}$ be the degree-D pseudoexpectation obtained from the degree-D matching pseudodistribution (Definition 4.4) for $(t+1)$-$\textsc{Pigeon}_N^M$. We show that $\widetilde{\mathbb{E}}$ is $(d, t, t-1)$-collision-free for $(t+1)$-$\textsc{Pigeon}_N^M$. By Theorem 4.7, this implies that $(t+1)$-$\textsc{Pigeon}_N^M$ does not depth-d reduce to t-$\textsc{Pigeon}_{N'}^{M'}$ for any M', N' with $M' \ge (t-1)N' + 1$.

Consider any t-witnessing family $\mathcal{F}$ for $(t+1)$-$\textsc{Pigeon}_N^M$ in which every conjunction has degree $\le d$. Without loss of generality, we can remove any conjunction C from $\mathcal{F}$ such that $\widetilde{\mathbb{E}}[C] = 0$. Thus we can assume that $\rho(C)$ encodes a partial matching for all $C \in \mathcal{F}$. Let $\mathcal{S} \subseteq \mathcal{F}$ be any collection of t conjunctions in $\mathcal{F}$, let $C_\mathcal{S} = \prod_{C\in\mathcal{S}} C$ denote the conjunction obtained by multiplying all conjunctions in $\mathcal{S}$.

Suppose that $C_\mathcal{S} \not\equiv 0$, and let $\rho = \rho(C_\mathcal{S})$ denote the corresponding partial assignment. Since each constituent conjunction of $C_\mathcal{S}$ is a partial matching, and there are only t conjunctions, it follows that ρ cannot witness a solution to $(t+1)$-$\textsc{Pigeon}_N^M$ since it can only contain a collision of at most t pigeons in any hole. A simple adversary strategy then implies that $DT((t+1)\text{-}\textsc{Pigeon}_N^M \restriction \rho) = \Omega(N)$, since we can respond to any unqueried pigeon by placing that pigeon into any hole with $\le t-1$ pigeons. Since the family is t-witnessing, it therefore follows that $C_\mathcal{S} \equiv 0$, i.e., $\rho(C_\mathcal{S})$ is an inconsistent partial assignment that tries to place at least one pigeon in two different holes. By Lemma 4.8, this means that $\widetilde{\mathbb{E}}[\mathcal{F}] \le t-1$, and therefore $\widetilde{\mathbb{E}}$ is a $(d, t, t-1)$-collision-free pseudoexpectation operator for $(t+1)$-$\textsc{Pigeon}_{N'}^{M'}$. Applying Theorem 4.7 completes the proof. □

Our main result separating the Pecking Order is now an immediate corollary of the previous theorem. We recall it here for convenience.

Theorem 1.8. If t is a constant and M, N, M', N' are parameters chosen so that $M' = \text{poly}(M)$, $M \geq tN + 1$, $M' \geq (t-1)N' + 1$, then $(t+1)$-PIGEON_N^M does not have an efficient black-box reduction to t-$\text{PIGEON}_{N'}^{M'}$. In particular, $(t+1)$-$\mathsf{PPP}^{dt} \not\subseteq t$-$\mathsf{PPP}^{dt}$ for any constant t, and so PiH^{dt} forms a strict hierarchy in the black-box setting.

Further, we note that we can prove a generalization of Theorem 4.9, separating the problem with parameter $a(n)$ from $b(n)$ whenever $a(n)$ is not polynomially close to $b(n)$. We prove this using the same technique as the theorem above, so we only provide a proof sketch here.

Theorem 1.11. Let $a(n), b(n)$ be functions such that $a(n)$ is larger than $b(n)$ and is not polynomially close to $b(n)$. Let M, N, M', N' be any parameters chosen so that $M' = \text{poly}(M)$, $M \geq (a(n) - 1)N + 1$, $M' \geq (b(n) - 1)N' + 1$. Then $a(n)$-PIGEON_N^M does not have an efficient black-box reduction to $b(n)$-$\text{PIGEON}_{N'}^{M'}$.

Proof Sketch. Similar to the proof of Theorem 4.9, we construct a $(d, b(n), b(n) - 1)$-Collision-free pseudoexpectation operator for $a(n)$-PIGEON_M^N using the Matching pseudodistribution (Definition 4.4) combined with Lemma 4.8. Here, we used that since $a(n)$ is not polynomially close to $b(n)$, no $b(n)$-collection of partial assignments can be witnessing for $a(n)$-PIGEON_N^M. Applying Theorem 4.7 completes the proof. □

Note that the statement of the above theorem separates all problems which do not have reductions guaranteed by Theorem 1.10. Combining the two, we can conclude Theorem 1.12.

Theorem 1.12. For any two functions $a(n), b(n)$, $a(n)$-$\mathsf{PPP}^{dt} = b(n)$-PPP^{dt} if and only if $a(n)$ and $b(n)$ are polynomially close.

5 Separations via the Compression Rate

In the previous we gave separations in the Pecking Order via collision number. In this section, we present another type of separation within the Pecking Order, which is due to the difference on the compression rate. Note that we are able to rule out *randomized* reductions in the following theorem.

Theorem 1.13. There is no randomized black-box reduction from PIGEON to n-PIGEON_N^M with $M \geq (n - 1 + c)N$ for a constant $c > 0$.

Corollary 5.1. $\mathsf{PPP}^{dt} \not\subseteq n$-$\mathsf{PWPP}^{dt}$.

Notation. We consider PIGEON instances with $N + 1$ pigeons and N holes. With a little abuse of notation, we denote a PIGEON instance by a string $h \in [N]^{N+1}$, where h_i is the hole that pigeon i gets mapped to. Assume that the n-PIGEON_N^M instance $f(h)$ reduced from h has $M(n')$ pigeons, $N(n')$ holes, and the solution is any n'-collision. We have $M' > N' \cdot (n' - 1 + c)$ for a constant $c > 0$ from the theorem statement. Formally, for any $i \in [M']$, pigeon i from $f(h)$ is mapped to hole $f_i(h)$. Without loss of generality, we assume all decision trees $(f_i, g_o)_{i \in [M'], o \in [M']^{n'}}$ have the same depth $d = \text{poly}(\log N)$.

We then proceed in four steps:

1. Finding an appropriate distribution D_N of hard PIGEON instances.

2. Showing that with high probability, there exists a "non-witnessing" solution in the n-$\textsc{Pigeon}_N^M$ instance $f(h)$ when h is drawn from D_N.
3. Arguing that the error probability is high if all decision trees (g_o) are depth-0.
4. Generalizing the error analysis to depth-d.

A hard distribution. By Yao's Minimax principle, it suffices to consider a family of distributions (D_N) of PIGEON instances, and then show that any deterministic low-depth black-box reduction (f_i, g_o) must be wrong with high probability. A natural choice of D_N is taking $h_1, \ldots, h_N$ to be a random permutation over $[N]$, while h_{N+1} is always set to 1. Now, the only possible solution is the index $i^* \in [N]$ such that $h_{i^*} = 1$. In the rest of this proof, we use i^* rather than the collision pair $(i^*, N+1)$ to indicate the solution of h; we also ignore h_{N+1} and assume h is permutation on $[N]$.

Find a non-witnessing solution. For any $i \in [M'], h \in D_N$, we say pigeon i in $f(h)$ is *non-witnessing* if the decision tree path in f_i realized by h is not witnessing, i.e., f_i does not query i^* when evaluating on h; otherwise, we call pigeon i *witnessing*.

Let $o = (i_1, \ldots, i_{n'})$ be any solution (i.e., an n'-collision) of $f(h)$. By taking the union of all decision tree paths in $(f_{i_j})_{j \in [n']}$ realized by h, we get a partial assignment h_o of size at most $d \cdot n' = \mathrm{poly}(\log N)$ in h. Without loss of generality, we assume this partial assignment h_o is also returned by the reduction as part of the solution.

We say o is a *non-witnessing* solution if for all $j \in [n']$, the pigeon i_j is non-witnessing, i.e., the partial assignment h_o does not witness the location of i^*; otherwise, o is a *witnessing* solution. Intuitively, a non-witnessing solution reveals almost no information about the key location i^*.

We now prove the following key lemma.

Lemma 5.2. When h is drawn from D_N, $f(h)$ has a non-witnessing solution with probability $1 - \mathrm{negl}(\log N)$.

Proof. Since h is a random permutation, and f_i only has d levels, for any pigeon $i \in [M']$ in $f(h)$, we have

$$\Pr_{h \sim D_N}[i \text{ is witnessing}] \leq \frac{d}{N} = \mathrm{negl}(\log N).$$

Using Markov's inequality, the probability that $f(h)$ has at least cN' witnessing pigeons is $\mathrm{negl}(\log N)$. In other words, with probability $1 - \mathrm{negl}(\log N)$, $f(h)$ has more than $(n'-1) \cdot N'$ non-witnessing pigeon.

Therefore, with probability $1 - \mathrm{negl}(\log N)$, $f(h)$ has a n'-collision with only non-witnessing pigeons, i.e., a non-witnessing solution.

□

Success probability for depth-0. We say a reduction is depth-k ($k \leq d$) if all the decision trees (g_o) are depth-k, while (f_i) are still depth-d. For a fixed family of depth-d decision trees (f_i), define p_k as the maximal success probability of any depth-k ($k \leq d$) reduction. We first consider the success probability of depth-0 reduction, i.e., $g_o \in [N]$ is a fixed location.

Let R, Q be the short-hand for PIGEON and n-$\textsc{Pigeon}_N^M$. Formally, our goal in this step is to show the following inequality.

Lemma 5.3.

$$p_0 \coloneqq \max_{f,g} \Pr_{h \sim D_N}[(h, g_o) \in R \Leftarrow (f(h), o) \in Q] < \mathrm{negl}(\log N).$$

Let $h^{(1)}$ be an instance of PIGEON with solution i^*. The key technical trick here is to roll a second dice, which helps us estimate the error probability. Specifically, we take a uniformly random index $j^* \in [N]$, and then generate $h^{(2)}$ from $h^{(1)}$ by swapping the solution from location i^* to j^*. Formally,

$$h^{(2)}_{j^*} = 1, h^{(2)}_{i^*} = h^{(1)}_{j^*};\ \ h^{(2)}_i = h^{(1)}_i, \forall i \in [N] \backslash \{i^*, j^*\}.$$

By symmetry, we know the marginal distribution of $h^{(2)}$ is also a random permutation, so we can calculate the success probability on $h^{(2)}$ instead.

Let $\mathsf{NonWit}(h, o)$ indicate the event that o is a non-witnessing solution for $f(h)$, and denote the uniform distribution over $[N]$ by U_N. We have the following observation regarding our second dice.

Lemma 5.4. For any $h^{(1)}$ and o,

$$\Pr_{j^* \sim U_N}[\mathsf{NonWit}(h^{(2)}, o) \mid \mathsf{NonWit}(h^{(1)}, o)] = 1 - \mathrm{negl}(\log N).$$

Proof. Since o is non-witnessing, the partial assignment h_o does not contain location i^*. Recall that $h^{(2)}$ is different to $h^{(1)}$ only on location i^* and j^*. Therefore, with $1 - \frac{|h_o|}{n} = 1 - \mathrm{negl}(\log N)$ probability on j^*, h_o is also a partial assignment of $h^{(2)}$, which implies that o is a non-witnessing solution of $f(h^{(2)})$. □

Now we are ready to prove Lemma 5.3.

Proof of Lemma 5.3. Fix an arbitrary depth-0 reduction (f, g). We give an overview before diving into the calculations. We first randomly draw $h^{(1)}$ from D_N, and there are two possible cases: either $f(h^{(1)})$ has a non-witnessing solution, or $f(h^{(1)})$ does not have any non-witnessing solution. The second case will happen with very low probability (Lemma 5.2), so we can assume $f(h^{(1)})$ has a non-witnessing solution o^*. We then roll a second dice $j \sim [N]$ and generate $h^{(2)}$ accordingly. We know that o^* is also a non-witnessing solution of $f(h^{(2)})$ with high probability (Lemma 5.4); among these $h^{(2)}$, the reduction could possibly be correct only when $j = g_{o^*}$.

Formally, let o^* be the first non-witnessing solution of $f(h^{(1)})$ in the lexicographical order, if $f(h^{(1)})$ has a non-witnessing solution; otherwise, let o^* to be the lexicographically first solution of $f(h^{(1)})$. We first consider whether $h^{(1)}$ has a non-witnessing solution.

$$\begin{aligned}
&\Pr_{h^{(2)} \sim D_N}[(h^{(2)}, g_o) \in R \Leftarrow (f(h^{(2)}), o) \in Q] \\
&= \Pr_{h^{(1)} \sim D_N, j^* \sim U_N}[(h^{(2)}, g_o) \in R \Leftarrow (f(h^{(2)}), o) \in Q] \\
&\leq \Pr_{h^{(1)}, j^*}[(h^{(2)}, g_o) \in R \Leftarrow (f(h^{(2)}), o) \in Q \mid \mathsf{NonWit}(h^{(1)}, o^*)] + \Pr_{h^{(1)}}[\neg \mathsf{NonWit}(h^{(1)}, o^*)] \qquad (2)
\end{aligned}$$

Note that the term $\Pr_{h^{(1)}}[\neg \mathsf{NonWit}(h^{(1)}, o^*)]$ is at most $\mathrm{negl}(\log N)$ from Lemma 5.2. So it remains to bound the first item in inequality (2), which is the success probability on $h^{(2)}$ condition on $h^{(1)}$ has a non-witnessing solution o^*.

$$
\begin{aligned}
&\Pr_{h^{(1)},j^*}[(h^{(2)}, g_o) \in R \Leftarrow (f(h^{(2)}), o) \in Q \mid \mathsf{NonWit}(h^{(1)}, o^*)] \\
&\leq \Pr_{h^{(1)},j^*}[(h^{(2)}, g_{o^*}) \in R \wedge \mathsf{NonWit}(h^{(2)}, o^*) \mid \mathsf{NonWit}(h^{(1)}, o^*)] \\
&\quad + \Pr_{h^{(1)},j^*}[\neg\mathsf{NonWit}(h^{(2)}, o^*) \mid \mathsf{NonWit}(h^{(1)}, o^*)] \\
&\leq \Pr_{h^{(1)},j^*}[(h^{(2)}, g_{o^*}) \in R \wedge \mathsf{NonWit}(h^{(2)}, o^*) \mid \mathsf{NonWit}(h^{(1)}, o^*)] + \mathrm{negl}(\log N) \qquad (3) \\
&= \Pr_{h^{(1)},j^*}[(j^* = g_{o^*}) \wedge \mathsf{NonWit}(h^{(2)}, o^*) \mid \mathsf{NonWit}(h^{(1)}, o^*)] + \mathrm{negl}(\log N) \\
&\leq \Pr_{h^{(1)},j^*}[(j^* = g_{o^*}) \mid \mathsf{NonWit}(h^{(1)}, o^*)] + \mathrm{negl}(\log N) \\
&= \frac{1}{N} + \mathrm{negl}(\log N) \\
&= \mathrm{negl}(\log N). \qquad (4)
\end{aligned}
$$

Note that inequality (3) comes from Lemma 5.4.

□

Success probability for depth-d. We now consider depth-k reduction for $k \leq d$, i.e., all (g_o) are depth-k. Define event $\mathsf{Hit\text{-}i}(h^{(1)}, o)$ to be true if o is a witnessing solution of $f(h^{(1)})$, or i^* is queried when evaluating g_o on $h^{(1)}$. Intuitively, $g_o(h^{(1)})$ has to *guess* a location if $\mathsf{Hit\text{-}i}(h^{(1)}, o)$ is not true,

The following lemma effectively reduces the problem to the case of depth-$(k-1)$.

Lemma 5.5. For a depth-k reduction (f_i, g_o),

$$
\Pr_{h^{(1)} \sim D_N}[\mathsf{Hit\text{-}i}(h^{(1)}, o) \Leftarrow (f(h^{(1)}), o) \in Q] \leq p_{k-1}
$$

Proof. We construct a depth-$(k-1)$ reduction (f_i, g'_o): For each possible $o \in [M']^{n'}$, the structure of g'_o is same as the first $k-1$ levels of g_o. If the solution i^* is witnessed by o itself or the first $k-1$ queries of g'_o, then g'_o will output i^*; otherwise, g'_o will output the location that is going be queried in g_o in the k-th level, if g_o is evaluated on the same input. Therefore, we have

$$
\Pr_{h^{(1)} \sim D_N}[\mathsf{Hit\text{-}i}(h^{(1)}, o) \Leftarrow (f(h^{(1)}), o) \in Q] = \Pr_{h^{(1)} \sim D_N}[(h^{(1)}, g'_o) \in R \Leftarrow (f(h^{(1)}), o) \in Q] \leq p_{k-1}.
$$

□

Define event $\mathsf{Hit\text{-}j}(h^{(1)}, o, j^*)$ to be true if the partial assignment h_o contains the location j^*, or j^* is queried when g_o is evaluated on $h^{(1)}$. Note that when both $\mathsf{Hit\text{-}i}(h^{(1)}, o)$ and $\mathsf{Hit\text{-}j}(h^{(1)}, o, j^*)$ are false, we have $g_o(h^{(1)}) = g_o(h^{(2)})$, because they only differ in location i^* and j^*. We have the following lemma using the same argument in Lemma 5.4.

Lemma 5.6. For any $h^{(1)}$ and o,

$$
\Pr_{j^* \sim U_N}[\neg\mathsf{Hit\text{-}j}(h^{(1)}, o, j^*) \mid \neg\mathsf{Hit\text{-}i}(h^{(1)}, o)] = 1 - \mathrm{negl}(\log N).
$$

Now, let us wrap everything up.

Proof of Theorem 1.13. We follow a similar strategy as in the proof of Lemma 5.3. For any depth-k reduction (f, g), there are two possibilities regarding to $h^{(1)}$: either $f(h^{(1)})$ has a (non-witnessing) solution o^* that $\mathsf{Hit\text{-}i}(h^{(1)}, o^*)$ is false, or $\mathsf{Hit\text{-}i}(h^{(1)}, o)$ is true for any solution o of $f(h^{(1)})$. The second case will happen with probability at most p_k by Lemma 5.5, and we also roll a second dice j^* to analyze the first case.

Formally, let o^* to be the lexicographically first solution of $f(h^{(1)})$ that $\mathsf{Hit\text{-}i}(h^{(1)}, o)$ is false, if such a solution exists; otherwise, let o^* to be the lexicographically first solution of $f(h^{(1)})$. We have

$$
\begin{aligned}
&\Pr_{h^{(1)}\sim D_N, j^*\sim U_N}[(h^{(2)}, g_o(h^{(2)})) \in R \Leftarrow (f(h^{(2)}), o) \in Q] \\
&\le \Pr_{h^{(1)}, j^*}[(h^{(2)}, g_o(h^{(2)})) \in R \Leftarrow (f(h^{(2)}), o) \in Q \mid \neg\mathsf{Hit\text{-}i}(h^{(1)}, o^*)] + \Pr_{h^{(1)}}[\mathsf{Hit\text{-}i}(h^{(1)}, o^*)] \\
&\le \Pr_{h^{(1)}, j^*}[(h^{(2)}, g_o(h^{(2)})) \in R \Leftarrow (f(h^{(2)}), o) \in Q \mid \neg\mathsf{Hit\text{-}i}(h^{(1)}, o^*)] + p_{k-1} \qquad (5) \\
&\le \Pr_{h^{(1)}, j^*}[(j^* = g_{o^*}(h^{(2)})) \wedge \neg\mathsf{Hit\text{-}j}(h^{(1)}, o^*, j^*) \mid \neg\mathsf{Hit\text{-}i}(h^{(1)}, o^*)] \\
&\quad + \Pr_{h^{(1)}, j^*}[\mathsf{Hit\text{-}j}(h^{(1)}, o^*, j^*) \mid \neg\mathsf{Hit\text{-}i}(h^{(1)}, o^*)] + p_{k-1} \\
&\le \Pr_{h^{(1)}, j^*}[(j^* = g_{o^*}(h^{(2)})) \wedge \neg\mathsf{Hit\text{-}j}(h^{(1)}, o^*, j^*) \mid \neg\mathsf{Hit\text{-}i}(h^{(1)}, o^*)] + p_{k-1} + \mathrm{negl}(\log N) \qquad (6) \\
&\le \Pr_{h^{(1)}, j^*}[(j^* = g_{o^*}(h^{(1)})) \mid \neg\mathsf{Hit\text{-}i}(h^{(1)}, o^*)] + p_{k-1} + \mathrm{negl}(\log N) \\
&= \frac{1}{N} + p_{k-1} + \mathrm{negl}(\log N) \\
&= p_{k-1} + \mathrm{negl}(\log N). \qquad (7)
\end{aligned}
$$

Inequality (5) uses Lemma 5.5, and inequality (6) comes from Lemma 5.6. From inequality (7), we have the success probability of any depth-d reduction

$$p_d \le d \cdot \mathrm{negl}(\log N) + p_0 = \mathrm{negl}(\log N). \qquad (8)$$

□

6 Relationship to Other Classes

6.1 Polynomial Local Search (PLS) and Polynomial Parity Argument (PPA)

Prior work in the literature has shown that $\mathsf{PPA}^{dt} \not\subseteq \mathsf{PPP}^{dt}$ [BCE+98] and $\mathsf{PLS}^{dt} \not\subseteq \mathsf{PPP}^{dt}$ [GHJ+22]. In this section, we prove that neither of these classes are contained in PAP^{dt}, recalled here.

Theorem 1.18. $\mathsf{PLS}^{dt} \not\subseteq \mathsf{PAP}^{dt}$ and $\mathsf{PPA}^{dt} \not\subseteq \mathsf{PAP}^{dt}$.

Both of these separations use the shared concept of *gluability*, which we first define.

Definition 6.1. Let C be a conjunction and let T be a decision tree. A *completion* of C by T is any conjunction of the form CC_ℓ where ℓ is any leaf of T.

Definition 6.2. A query total search problem $R \subseteq \{0,1\}^n \times O$ is (d, t)-*gluable* if for every conjunction C of degree at most d there is a depth $O(d)$ decision tree T_C such that the following holds. Let $C_1, C_2, \ldots, C_t$ be any sequence of t consistent conjunctions, and let $C'_1, \ldots, C'_t$ be any sequence of t conjunctions chosen so that C'_i is a completion of C_i by T_{C_i}. If C'_i is non-witnessing for each $i = 1, 2, \ldots, t$ and $C'_1 C'_2 \cdots C'_t$ is consistent, then $DT(R \restriction \rho(C'_1 C'_2 \cdots C'_t)) \ge d$.

A weaker notion of gluability was introduced by Göös et al. [GHJ+22] in the case where $t = 2$, in order to prove $\mathsf{PLS}^{dt} \not\subseteq \mathsf{PPP}^{dt}$, although the idea also implicitly appears in [BCE+98]. We show a new result for gluability related to collision-freedom: if a problem R is gluable *and* it admits a non-witnessing pseudoexpectation operator $\widetilde{\mathbb{E}}$, then one can show automatically that $\widetilde{\mathbb{E}}$ is also collision-free, and Theorem 4.7 could be applied. We prove this now in a strong form.

Definition 6.3. Let $R \subseteq \{0,1\}^n \times O$ be a query total search problem, and let $\widetilde{\mathbb{E}}$ be a degree-d pseudoexpectation operator. Then $\widetilde{\mathbb{E}}$ is *ε-nonwitnessing for R* if the following property holds:

- **ε-Nonwitnessing.** $\widetilde{\mathbb{E}}[C] \leq \varepsilon$ for any degree-d conjunction C witnessing R.

Note that 0-nonwitnessing is synonymous with R-nonwitnessing (cf. Definition 4.1).

Our goal now is to prove the following theorem.

Theorem 6.4. Let $R = \{R_n \subseteq \{0,1\}^n \times O_n\}$ be a query total search problem not contained in PPADS^{dt}, and suppose $t = O(\mathrm{poly}(\log n))$. If R is $(p(\log(n)), t)$-gluable for any polynomial function p, then $R \notin t\text{-}\mathsf{PPP}^{dt}$.

To prove this theorem, we will need to define the *Sherali-Adams* proof system. To avoid introducing proof-complexity preliminaries, we will state the definition of the proof system directly in terms of the total search problems R, although we refer the interested reader to [GHJ+22] for technical details. If $R \subseteq \{0,1\}^n \times O$ is a query total search problem in TFNP, then we define a related unsatisfiable CNF formula $\neg\mathrm{Total}(R)$ that encodes the (false) statement "R is not total". Formally,

$$\neg\mathrm{Total}(R) := \bigwedge_{o \in O} \bigwedge_{\ell \in L_1(T_o)} \neg C_\ell$$

where $\{T_o\}_{o \in O}$ is the family of $\mathrm{poly}(\log(n))$-depth decision trees witnessing solutions of R, and $L_1(T_o)$ is the set of 1-leaves of the decision tree T_o. We will be interested in *refutations* of the formula $\neg\mathrm{Total}(R)$ — in other words, proofs of the tautology "R is total". Recall that a *conical junta* is any non-negative linear combination of conjunctions, that is an expression of the form $J = \sum_D \lambda_D D$, where each D is a conjunction and $\lambda_D > 0$. The *degree* of J is the maximum degree of any conjunction D appearing in the sum. The *magnitude* of J, denoted $||J||$, is $\max_D \lambda_D$.

Definition 6.5. Let $R \in \mathsf{TFNP}^{dt}$ be a total query search problem. A *$\mathbb{Z}$-Sherali-Adams proof of totality for R* is a sequence of conical juntas $\Pi = (J, J_\ell)_{o \in O, \ell \in L_1(T_o)}$ over the variables of R such that

$$-1 = \sum_{o \in O} \sum_{\ell \in L_1(T_o)} -J_\ell C_\ell + J$$

where we are working in multilinear polynomial algebra. The *degree* of the refutation Π is $\deg(\Pi) := \max\{\deg J_o C_o\}_o \cup \{\deg J\}$. The *magnitude* of the refutation Π is $||\Pi|| := \max\{||J_o||\}_o \cup \{J\}$.

Göös et al. [GHJ+22] proved that that $R \in \mathsf{PPADS}^{dt}$ if and only if it admits Sherali-Adams proofs with low degree and magnitude.

Theorem 6.6 ([GHJ+22])**.** Let $R \in \mathsf{TFNP}^{dt}$ be a total query search problem. Then $R \in \mathsf{PPADS}^{dt}$ if and only if for some constant c, R admits a $\log^c(n)$-degree, $n^{\log^c(n)}$-magnitude $\mathbb{Z}$-Sherali-Adams proof of totality.

One can show that *any* lower bound against degree-d Sherali-Adams with small magnitude implies the existence of a weak pseudoexpectation operator (in other words, ε-nonwitnessing operators are *complete* for unary Sherali-Adams lower bounds). The proof of this follows the usual proof of completeness for Sherali-Adams proofs via convex duality (see e.g. [FKP19]), and was first observed by Hubáček, Khaniki, and Thapen [HKT24].

Theorem 6.7 ([HKT24])**.** If there is no degree-d $\mathbb{Z}$-Sherali-Adams proof of totality for R of magnitude $\leq k$, then there is a degree-d, $1/k$-nonwitnessing pseudoexpectation operator for R.

We can now prove Theorem 6.4.

Proof of Theorem 6.4. Let $R \subseteq \{0,1\}^n \times O$ be any $(p(\log n), t)$-gluable total search problem not contained in PPADS. Suppose by way of contradiction that $R \in t\text{-PPP}^{dt}$, and so R has a depth-d reduction to $t\text{-PIGEON}_N^{(t-1)N+1}$ where $d = \log^{C_0} n$ for some universal constant C_0 and $N = n^{\log^{C_0} n}$. We use the fact that R is (d^*, t)-gluable for $d^* = \log^{2C_0} n = \omega(d)$. Since $R \notin \text{PPADS}^{dt}$, it follows from Theorem 6.6 that for *every* positive constant α, there is no degree-$\log^\alpha n$, $\mathbb{Z}$-Sherali-Adams proof of totality for R with magnitude $n^{\log^\alpha n}$. By Theorem 6.7, the non-existence of a $\mathbb{Z}$-Sherali-Adams proof implies that there is a degree-$\log^\alpha n$, $1/n^{\log^\alpha n}$-nonwitnessing pseudoexpectation operator $\widetilde{\mathbb{E}}$ for R. With this in mind, let $\widetilde{\mathbb{E}}$ be a degree-$t\log^{4C_0} n$, ε-nonwitnessing pseudoexpectation operator for R with $\varepsilon = 1/n^{t\log^{4C_0} n}$. Our goal is to show that $\widetilde{\mathbb{E}}$ is $(d, t, t - 1 + o(1/N))$-collision-free, which will contradict the existence of the assumed reduction to t-PIGEON by Theorem 4.7.

Let $\mathcal{F}$ be any t-witnessing family of width $\leq d$ conjunctions, and let $\mathcal{C}(\mathcal{F})$ be the family obtained by replacing each conjunction C in $\mathcal{F}$ with all of its completions CC_ℓ for $\ell \in L(T_C)$ guaranteed by (d^*, t)-gluability. We first observe that $\widetilde{\mathbb{E}}[\mathcal{C}(\mathcal{F})] = \widetilde{\mathbb{E}}[\mathcal{F}]$, noting that the degree of $\widetilde{\mathbb{E}}$ is $\omega(d)$ and therefore both of these expressions are well-defined. To see this, consider any conjunction $C \in \mathcal{F}$. If T_C is the decision tree completing C in the definition of gluability, then an easy induction on the depth of the decision tree shows that

$$1 = \sum_{\ell \in L(T_C)} C_\ell$$

where the equality is between multilinear polynomials. This implies that $C = \sum_{\ell \in L(T_C)} CC_\ell$ and thus

$$\widetilde{\mathbb{E}}[C] = \sum_{\ell \in L(T_C)} \widetilde{\mathbb{E}}[CC_\ell].$$

It immediately follows that $\widetilde{\mathbb{E}}[\mathcal{F}] = \widetilde{\mathbb{E}}[\mathcal{C}(F)]$, since $\mathcal{C}(\mathcal{F})$ is obtained by replacing each conjunction C with its completions. This means that it suffices to bound the weight of $\mathcal{C}(\mathcal{F})$ instead of $\mathcal{F}$. Let $W \subseteq \mathcal{C}(\mathcal{F})$ be the collection of all conjunctions in $\mathcal{C}(\mathcal{F})$ that are themselves witnessing, and let $X = \mathcal{C}(\mathcal{F}) \setminus W$ be the remaining conjunctions, and note that $\widetilde{\mathbb{E}}[\mathcal{C}(\mathcal{F})] = \widetilde{\mathbb{E}}[X] + \widetilde{\mathbb{E}}[W]$.

First, let's bound the weight of $\widetilde{\mathbb{E}}[W]$. Since every conjunction $C' \in W$ is witnessing, it follows that $\widetilde{\mathbb{E}}[C'] \leq \varepsilon$ since $\widetilde{\mathbb{E}}$ is ε-nonwitnessing for R. This means that $\widetilde{\mathbb{E}}[W] \leq |W|\varepsilon \leq \varepsilon n^{O(d)}$, since every conjunction in W has width at most $O(d)$.

We now bound the weight of X. Let $C_1, C_2, \ldots, C_t$ be any sequence of distinct conjunctions chosen from X. Since $\mathcal{F}$ is t-witnessing, $\mathcal{C}(\mathcal{F})$ is also t-witnessing, and this implies that if $C' = C_1 C_2 \ldots C_t$ is consistent, then $DT(R \restriction \rho(C')) = O(d)$. However, by the definition of (d^*, t)-gluability, whenever C' is consistent and each C_i is non-witnessing, then $DT(R \restriction \rho(C')) \geq d^*$. Since $d^* = \omega(d)$, these two facts together imply that if $C_1, C_2, \ldots, C_t \in X$ are distinct conjunctions, then $C_1 C_2 \cdots C_t$ must be inconsistent, since no individual conjunction X can be witnessing. Therefore, all such C' composed of conjunctions from X are inconsistent, and so applying Lemma 4.8 we observe that $\widetilde{\mathbb{E}}[X] \leq t - 1$.

Combining the two weight bounds yields $\widetilde{\mathbb{E}}[\mathcal{F}] = \widetilde{\mathbb{E}}[\mathcal{C}(\mathcal{F})] = \widetilde{\mathbb{E}}[W] + \widetilde{\mathbb{E}}[X] = t - 1 + \varepsilon n^{O(d)}$. Our choice of parameters then implies that

$$\varepsilon n^{O(d)} = \frac{n^{O(\log^{C_0} n)}}{n^{t \log^{4C_0} n}} = o\left(\frac{1}{N}\right),$$

where the last equality follows since $N = n^{\log^{C_0} n}$. This means that the pseudoexpectation $\widetilde{\mathbb{E}}$ for R is a $(d, t, t - 1 + o(1/N))$-collision-free pseudoexpectation operator. This is a contradiction to Theorem 4.7, and it follows that $R \notin t\text{-PPP}^{dt}$. □

In the remainder of this section, we prove the required gluability results for PPA and PLS.

PPA^{dt} is gluable. We first introduce the defining problem for the class PPA^{dt}, called LEAF.

Definition 6.8 (LEAF$_n$)**.** This problem is defined on a set of n nodes, denoted by $[n]$, where the node 1 is "distinguished". For input, we are given a neighbourhood $N_u \subseteq [n]$ of size $|N(u)| \leq 2$ for each node $u \in [n]$. Given this list of neighbourhoods, we create an undirected graph G where we add an edge uv if and only if $u \in N(v)$ and $v \in N(u)$. We say $u \in [n]$ is a *leaf* if it has in-degree 1 and out-degree 0. The goal of the search problem is to output either

1. 1, if 1 is not a leaf in G, or *(no distinguished leaf)*
2. $u \neq 1$, if u is a leaf in G. *(proper leaf)*

The class PPA^{dt} contains all query total search problems with poly(log(n))-complexity reductions to LEAF.

The seminal work by Beame et al. [BCE+98] proved that LEAF $\notin$ PPP^{dt}, which implies that LEAF $\notin$ PPADS^{dt}. Therefore, by Theorem 6.4, to prove PPA$^{dt} \not\subseteq$ PAP^{dt}, we need to show that LEAF$_n$ is (poly(log n), poly(log n))-gluable.

Lemma 6.9. LEAF$_n$ is $(p(\log(n)), p(\log(n)))$-gluable for any polynomial p.

Proof. Let $d, t = \text{poly}(\log(n))$. Let C be any conjunction of degree d over the variables of LEAF$_n$. The decision tree T_C completing C does the following: if C ever queries a node $u \in [n]$, receiving a neighbourhood $N(u)$, T_C queries the nodes in $N(u)$ as well. In total, this requires $O(d)$ more queries. Now, let $C_1, C_2, \ldots, C_t$ be any sequence of t consistent conjunctions, and let $C'_1, C'_2, \ldots, C'_t$ be any sequence of t conjunctions chosen so that C'_i is a completion of C_i by T_{C_i}. Assume that these completions are all consistent and non-witnessing, and let $C' = C'_1 C'_2 \cdots C'_t$. First suppose by way of contradiction that C' witnesses a LEAF$_n$ solution u. But in either case, the fact that u is a solution must have been witnessed by some C'_i in the sequence, since we have explicitly queried all the neighbourhoods of nodes in each C_i. This is a contradiction, and thus C' must be non-witnessing.

Let us now show LEAF$_n \restriction \rho(C')$ has large decision tree depth. We can give a simple adversary argument as follows. Let U be the set of all nodes currently queried — initially, U contains all nodes queried by C', and thus $|U| = \text{poly}(\log(n))$. Say a node u is on the *boundary* of U if it has not been queried, but appears as a neighbour of a queried node. Consider a decision tree T querying nodes in LEAF$_n$. If T queries a node v not in the boundary of U, then output $N(v) = \emptyset$. Otherwise, suppose T queries a node v in the boundary of U. If two nodes $\{u_1, u_2\} \subseteq U$ have v in their neighbour set, then set $N(v) = \{u_1, u_2\}$. Otherwise, if just one node $u \in U$ has v in its neighbour set, then set $N(v) = \{u, w\}$, where w is any node not in U nor the boundary of U. Since the degree of every node is ≤ 2 and $|U| = \text{poly}(\log(n))$, we can clearly continue this adversary strategy for $\Omega(n)$ queries. Therefore $DT(\text{LEAF}_n \restriction \rho(C')) = \Omega(n)$, and thus LEAF$_n$ is $(\text{poly}(\log(n)), \text{poly}(\log(n)))$-gluable. □

Corollary 6.10. PPA$^{dt} \not\subseteq$ PAP^{dt}.

PLS^{dt} is gluable. Let us first recall the Sink-of-Dag problem, also denoted SoD, which is the defining problem for PLS. Our definition follows that of Göös et al. [GHJ+22].

Definition 6.11 (Sink-of-Dag)**.** The SoD_n problem is defined on the $[n] \times [n]$ grid, where the node $(1,1)$ is "distinguished". As input, for each grid node $u = (i,j) \in [n] \times [n]$, we are given an "active bit" $a_u \in \{0,1\}$. Furhter, we are given a *successor* $s_u \in [n]$, interpreted as naming a node $(i+1, s_u)$ on the next row. We say a node u is *active* if $a_u = 1$, otherwise it is *inactive*. A node u is a *proper sink* if u is inactive but some active node has u as a successor. The goal of the search problem is to output either

1. $(1,1)$, if $(1,1)$ is inactive *(inactive distinguished source)*
2. (n,j), if (n,j) is active, *(active sink)*
3. (i,j) for $i \leq n-1$, if (i,j) is active and its successor is a proper sink. *(proper sink)*

We define the class PLS^{dt} to be all total query search problems with $\mathrm{poly}(\log(n))$-complexity Sink-of-Dag-formulations.

One of the main results of Göös et al. is the following.

Theorem 6.12 ([GHJ+22], Corollary 1.)**.** $\mathsf{PLS}^{dt} \not\subseteq \mathsf{PPADS}^{dt}$.

Göös et al. already showed that the SoD_n problem is $(\mathrm{poly}(\log(n)), 2)$-gluable, but it is not hard to see that their proof generalizes to show that it is actually $(\mathrm{poly}(\log(n)), t)$-gluable for any $t = \mathrm{poly}(\log n)$. The proof of the following lemma closely follows the argument of [GHJ+22, Lemma 13].

Lemma 6.13. SoD is $(p(\log(n)), p(\log(n)))$-gluable for any polynomial p.

Proof. Let C be any conjunction of degree $d = \mathrm{poly}(\log(n))$, defined over the variables of SoD_n. The decision tree T_C completing C starts by checking whether C queries any active node below row $n-d-2$. If yes, T picks any one such active node and follows the successor path until a sink is found, making the completion witnessing. Note that this step incurs at most $O(d)$ queries. Finally, T ensures that any query to a successor variable in C is followed by a query to the active bit of the successor. This also costs $O(d)$ further queries. Thus the depth of T_C is $O(d) = \mathrm{poly}(\log(n))$.

As in the definition of gluability, let $C_1, C_2, \ldots, C_t$ be any sequence of t consistent conjunctions, and let $C'_1, C'_2, \ldots, C'_t$ be any sequence of completions of those conjunctions by their respective decision trees. Suppose that C'_i is non-witnessing for each i and that $C'_1, C'_2, \ldots, C'_t$ are consistent, and suppose by way of contradiction that the conjunction $C' = C'_1 C'_2 \cdots C'_t$ is witnessing. If it reveals a SoD solution u of type (1) or (2), then it must be that for some i, C'_i queries an active bit of u: a contradiction with the fact that C'_i is non-witnessing. On the other hand, if C' reveals a solution u of type (3), then it must be that for some i, C'_i checks for the successor s_u of u, but the completion T_{C_i} forces this check to be followed by a query to the active bit of s_u, making one of the initial partial assignments witnessing as well. Hence C' is non-witnessing.

We finally argue that $R_n \restriction \rho(C')$ has query complexity greater than d by describing an adversary that can fool any further ℓ queries to p without witnessing a solution. Recall that p makes no queries to nodes below row $n-d-2$. The adversary answers queries as follows. If the successor pointer of an active node is queried, then we answer with a pointer to any unqueried node on the next row and make it active (there always exists one as $d \ll n$). If a node u is queried that is not the successor of any node, we make u inactive ($a_u = 0$ and s_u is arbitrary). This ensures that a solution can only lie on the very last row n, which is not reachable in ℓ queries starting from row $n-d-2$. □

We can thus infer Theorem 1.18, which we restate here for convenience.

Theorem 1.18. $\mathsf{PLS}^{dt} \not\subseteq \mathsf{PAP}^{dt}$ and $\mathsf{PPA}^{dt} \not\subseteq \mathsf{PAP}^{dt}$.

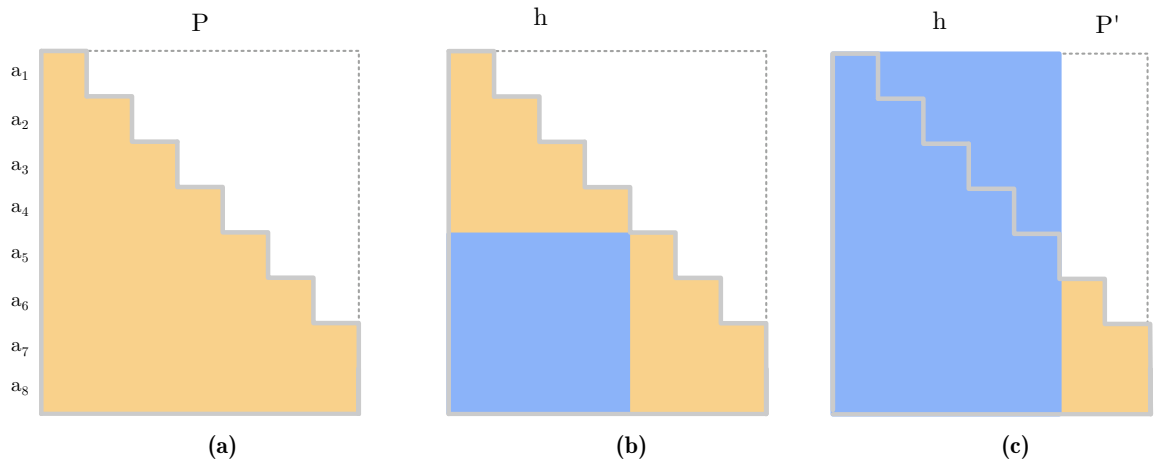

Figure 2: UPLC solutions and reductions. **(a)** illustrates the lower-triangular condition of solutions to UPLC. **(b)** is an illustration of Lemma 6.15. The blue shaded region is a valid solution to $n/2$-PWPP which is a subset of the yellow shaded region, the lower-triangular collision returned by UPLC. **(c)** is an illustration of Lemma 6.16. The blue shaded region is the collision returned by n-PIGEON, and the yellow shaded region is computed by solving a $\log n$ size UPLC instance by brute force.

6.2 Iterated Pigeonhole

In this section, we discuss the relationship between the Pecking Order and classes PLC and UPLC defined by [PPY23]. We give an alternate definition of UPLC which is more convenient for our reductions. The class UPLC is defined by its complete problem with the same name UPLC, specified in below.

Definition 6.14 (UPLC)**.** A universe of elements $\{0,1\}^n$ is considered.

Input A circuit $P : \{0,1\}^n \mapsto \{0,1\}^{n-1}$.

Solutions A set of n distinct strings $a_1 \dots a_n$ such that for every j the strings $P(a_j), P(a_{j+1}) \dots P(a_n)$ agree on the prefix of length j.

We get the above definition of UPLC by concatenating the circuits in the original definition [PPY23]. We refer to the type of solution reported by UPLC as the "lower-triangular condition". This is illustrated in Figure 2a.

Lemma 6.15. $n/2$-$\text{PIGEON}^N_{\sqrt{N}}$ is in UPLC.

Proof. Given an instance (n, h) of $n/2$-$\text{PIGEON}^N_{\sqrt{N}}$, where $h : \{0,1\}^n \mapsto \{0,1\}^{n/2}$, we construct an instance P of UPLC by simply appending 0s to the circuit.

$$P(x_1 x_2 \dots x_{n-1}) = h(x_1 \dots x_{n/2}) 0^{n/2-1}.$$

The solution to UPLC gives us the lower triangular condition for some set of strings $\{a_0, a_1 \dots a_n\}$ which in particular gives us an $n/2$-collisions for h, since $a_{\frac{n}{2}} \dots a_n$ agree on the first $n/2$ bits. This proof is illustrated in Figure 2b. □

Theorem 1.16. $\text{UPLC}^{dt} \not\subseteq \text{SAP}^{dt}$.

Proof. This is by combining Lemma 6.15 with Theorem 4.9. □

Further, we complete a fine-grained understanding of the position of UPLC in the Pecking Order by placing it in n-PWPP.

Lemma 6.16. UPLC $\subseteq$ n-PWPP.

Proof. Without loss of generality, we assume n is a power of 2. Given the circuit $P : \{0,1\}^n \mapsto \{0,1\}^{n-1}$ of a UPLC instance, we "split" the output into two parts: the first $n - \log n$ bits are considered as an instance of n-PIGEON$^N_{N/n}$ and the remaining $\log n - 1$ bits are considered as an instance of UPLC in a $\log n$ scale.

We first solve the n-PIGEON$^N_{N/n}$ instance defined by a mapping $h : \{0,1\}^n \mapsto \{0,1\}^{n-\log n}$, where h outputs the first $n - \log n$ bits of P. Note that n-PIGEON$^N_{N/n} \in n$-PWPP. Suppose we get an n-collision $(a_1, \ldots, a_n)$ of h in the first step.

We then consider the UPLC instance P' on the universe $U = \{a_1, \ldots, a_n\}$, where $P' : U \mapsto \{0,1\}^{\log n - 1}$ is defined by the remaining $\log n - 1$ bits of P. We can solve this much smaller UPLC instance in polynomial time, and get a set of solution $(a_{i_1}, \ldots, a_{i_{\log n}})$.

Finally, we get the solution to the original UPLC instance by rearranging the elements in U: we put $(a_{i_1}, \ldots, a_{i_{\log n}})$ in the end with the same order, and put everything else in before with arbitrary order. The first $n - \log(n)$ columns of the lower-triangular condition are satisfied by finding the n-collision of h, and the remaining $\log n - 1$ columns are fulfilled by solving the small UPLC instance P'.

This proof is illustrated in Figure 2c. □

Combining Theorem 1.13, Lemma 6.16, and the fact that PPP $\subseteq$ PLC [PPY23], we get two more black-box separations regarding to UPLC.

Theorem 1.17. PPP$^{dt} \not\subseteq$ UPLC^{dt}. Consequently, PLC$^{dt} \not\subseteq$ UPLC^{dt}.

A stronger version of UPLC. We observe that UPLC is still total if the universe has size $2^{n-1}+1$, and we can even ask for more matching bits between those pigeons. To see this, we apply the *non-adaptive* version of the iterative PHP: Start with the universe $U_0 = [2^{n-1}+1]$. In the i-th step, the i-th bit of P divides the current set U_{i-1} into two subsets and sets U_i to be the larger one (break ties arbitrarily). In this way, for any $i \in [n-1]$, all the elements in U_i agree on the first i bits of P, and $|U_{n-i}| \geq 2^i$. Therefore, we can pick

1. two elements (a_{n-1} and a_n) from U_{n-1} that match on all $n-1$ bits of P,
2. at least 2 more elements from U_{n-2} that match on first $n-2$ bits with a_n, ...,
3. and in general, at least 2^{i-1} more elements from U_{n-i} that match on first $n-i$ bits with a_n.

We now define a new problem T-UPLC ("T" for tight) to study the hardest computational problem corresponding to the non-adaptive iterative pigeonhole principle described above.

Definition 6.17 (T-UPLC)**.** Let $n = 2^\ell$ be a power of 2, and a universe of elements $U = [2^n + 1]$ is considered.

Input A circuit $P : U \mapsto \{0,1\}^n$.

Solutions A set of $n+1$ distinct strings $a_0 \ldots a_n$ such that for every $i \in \{0, 1, \ldots, \ell\}$, the strings $P(a_{n-2^i}), P(a_{n-2^i+1}) \ldots P(a_n)$ agree on the prefix of length $n-i$.

We show the counter-intuitive result that the non-adaptive iterative pigeonhole principle is equivalent to the generalized pigeonhole principle computationally.

Lemma 6.18. T-UPLC is PAP-complete.

Proof Sketch. Let $n = 2^\ell$. Note that the solution $\{a_0 \dots a_n\}$ for a T-UPLC instance (n, P) is also the solution for the $(n+1)$-PIGEON$^{N+1}_{N/n}$ instance (n, h) with h being the first $n - \ell$ bits of P. Therefore, T-UPLC is PAP-hard.

To show that T-UPLC is also in PAP, we use the same two-stage argument as in Lemma 6.16: The first $n - \ell$ bits of P are considered as an instance of $(n+1)$-PIGEON$^{N+1}_{N/n}$ and we get an $(n+1)$-collision $U' \coloneqq \{a_1, \dots, a_{n+1}\}$ in this step. The remaining ℓ bits of P are interpreted as an instance of T-UPLC in a $\log n$ scale with the universe U', each can be solved trivially. □

6.3 Total Function BQP

In this section, we put the problem defined by Yamakawa-Zhandry [YZ22] for their breakthrough result in the Pecking Order. Yamakawa-Zhandry's problem was first defined relative to a random oracle. It was later adapted to constitute a total problem ([YZ22], Section 6), which is the version we considered in this paper.

Theorem 1.19. Yamakawa-Zhandry's Problem is contained in PAP.

We start with formally defining the Yamakawa-Zhandry's problem.

Definition 6.19 (Yamakawa-Zhandry's Problem [YZ22], Simplified)**.** Fix an error correcting code $C \subseteq \Sigma^n$ on alphabet Σ. Let (h_k) be a family of λ-wise independent functions from C to $\{0,1\}^n$.

Input The input encodes n mapping $f_1, \dots, f_n : \Sigma \to \{0,1\}$. We define $f : \Sigma^n \to \{0,1\}^n$ as $f(a_1 a_2 \dots a_n) \coloneqq (f_1(a_1), f_2(a_2), \dots, f_n(a_n))$.

Solutions The goal is to find a key k and t codewords $c^{(1)} \dots, c^{(t)} \in C$ such that $f(c^{(i)}) \oplus h_k(c^{(i)}) = 0^n$, for all $i \in [t]$.

The parameters satisfy $n < \lambda \ll t = \text{poly}(n)$, $|C| \geq 2^{2n}$, and $|\Sigma| = 2^{\text{poly}(n)}$.

This definition is slightly different to the one appeared in [YZ22][7], which is simpler to analyze. A *folded Reed-Solomon* code (see [GRS23], Section 17.1) with a certain parameter setting is chosen for the ECC C in [YZ22]. The structure of the ECC is crucial for the exponential quantum speed-up. It is straightforward to verify that the problem as stated here preserves the quantum upper bound and the classical lower bound. We assume the λ-wise independent functions family (h_k) is implemented by the well-known low-degree polynomial construction.

Definition 6.20 ([Vad12], Section 3.5.5)**.** Let $\mathbb{F}$ be a finite field. Define the family of functions $\mathcal{H} = \{h_{a_0, a_1 \dots a_{\lambda-1}} : \mathbb{F} \mapsto \mathbb{F}\}$, where each $h_{a_0, a_1 \dots a_{\lambda-1}} = a_0 + a_1 x + a_2 x^2 \dots a_{\lambda-1} x^{\lambda-1}$ for $a_0, a_1 \dots a_{\lambda-1} \in \mathbb{F}$.

It is noted in [Vad12] that this family forms a λ-wise independent set.

A further simplification. To show Theorem 1.19, we consider a simplification of the Yamakawa-Zhandry's Problem, in which we assume all the codewords have n distinct letters, and no letter appears in two different codewords. We call this new problem ALLZEROCOLUMN. The rationale behind the naming is that we can now imagine we are given an *arbitrary* 0/1 matrix F of size $n \times |C|$ representing the mapping f, where the i-th column of F corresponds to the result of applying f to the i-th codeword.

[7] In [YZ22], the input contains t mappings $f_1, \dots, f_t$, and the goal is find a key k and t codewords $c^{(1)} \dots, c^{(t)} \in C$ such that $f_i(c^{(i)}) \oplus h_k(c^{(i)}) = 0^n, \forall i \in [t]$.

Definition 6.21 (AllZeroColumn)**.** Fix a 0-1 matrix family $\mathcal{H}$ of size $n \times m$, where the entries are λ-wise independent, implemented by the low-degree polynomial construction.

Input The input is a 0-1 matrix F of size $n \times m$.

Solutions The goal is to find t indices $j_1, \ldots, j_t \in [m]$ and a matrix $H_k \in \mathcal{H}$ (which could be succinctly represented by a key k), such that for any $i \in [t]$, the j_i-th column of matrix $F \oplus H_k$ is $\mathbf{0}^n$.

The parameters satisfy $n < \lambda \ll t = \text{poly}(n)$, and $2^n \cdot t \leq m = 2^{\text{poly}(n)}$.

As such, the Yamakawa-Zhandry's problem can be seen as AllZeroColumn with a *promise*: F but has the structure imposed by the ECC C in the definition of the Yamakawa-Zhandry's problem. Theorem 1.19 is then implied by the following lemma.

Lemma 6.22. AllZeroColumn is contained in PAP.

Proof. Given an AllZeroColumn instance specified by a 6-tuple $(n, m, \lambda, t, F, \mathcal{H})$, we construct the following t-Pigeon$_N^M$ instance for $M = m \cdot \frac{|\mathcal{H}|}{2^n}$ and $N = |\mathcal{H}|$.

Pigeons Each pigeon is a pair (j, k). Here j is in $[m]$, and k is a key for the family $\mathcal{H}$ such that the j-th column of matrix $F \oplus H_k$ is $\mathbf{0}^n$.

Holes Each hole is a key k.

Pigeon (j, k) is now mapped to hole k. Recall that we have $n < \lambda$, thus, each index j could pair with exactly 2^{-n} fractions of keys to form a valid pigeon. Hence, there are $M = m \cdot \frac{|\mathcal{H}|}{2^n}$ pigeons. Since $m \geq 2^n \cdot t$ in the AllZeroColumn instance, we can verify that $M \geq t \cdot N$. Also, any t-collision in the t-Pigeon$_N^M$ instance directly corresponds to a solution of the AllZeroColumn instance.

Since we implemented $\mathcal{H}$ using the low-degree polynomial construction Definition 6.20, this reduction can be implemented in $\text{poly}(n)$ time in the white-box setting using polynomial interpolation. □

Acknowledgements

We thank Beach House for providing topical procrastination music [Hou15]. We thank Mika Göös, Scott Aaronson, David Zuckerman, Alexandors Hollender, Gilbert Maystre, Ninad Rajgopal, Chetan Kamath for discussions about the Pigeonhole principle. We also thank Scott Aaronson for suggesting the name Pecking Order. SJ thanks Mika Göös for countless insightful discussions about TFNP and complexity theory more generally.

SJ and JL are supported by Scott Aaronson's Vannevar Bush Fellowship from the US Department of Defense, the Berkeley NSF-QLCI CIQC Center, a Simons Investigator Award, and the Simons "It from Qubit" collaboration. ZX is supported by NSF award CCF-2008868 and the NSF AI Institute for Foundations of Machine Learning (IFML).

References

[AA14] Scott Aaronson and Andris Ambainis. The need for structure in quantum speedups. *Theory Comput.*, 10:133–166, 2014. doi:10.4086/toc.2014.v010a006.

[AS04] Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. *J. ACM*, 51(4):595–605, 2004. doi:10.1145/1008731.1008735.

[BB14] Arnold Beckmann and Samuel R. Buss. Improved witnessing and local improvement principles for second-order bounded arithmetic. *ACM Trans. Comput. Log.*, 15(1):Art. 2, 35, 2014. doi:10.1145/2559950.

[BCE+98] Paul Beame, Stephen Cook, Jeff Edmonds, Russell Impagliazzo, and Toniann Pitassi. The relative complexity of NP search problems. volume 57, pages 3–19. 1998. 27th Annual ACM Symposium on the Theory of Computing (STOC'95) (Las Vegas, NV). doi:10.1006/jcss.1998.1575.

[BDRV18] Itay Berman, Akshay Degwekar, Ron D. Rothblum, and Prashant Nalini Vasudevan. Multi-collision resistant hash functions and their applications. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, volume 10821 of *Lecture Notes in Computer Science*, pages 133–161. Springer, 2018. doi:10.1007/978-3-319-78375-8_5.

[BFH+23] Romain Bourneuf, Lukáš Folwarczný, Pavel Hubáček, Alon Rosen, and Nikolaj I. Schwartzbach. PPP-completeness and extremal combinatorics. In *14th Innovations in Theoretical Computer Science Conference*, volume 251 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 22, 20. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2023. doi:10.4230/lipics.itcs.2023.22.

[BFI23] Sam Buss, Noah Fleming, and Russell Impagliazzo. TFNP characterizations of proof systems and monotone circuits. In *14th Innovations in Theoretical Computer Science Conference*, volume 251 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 30, 40. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2023. doi:10.4230/lipics.itcs.2023.30.

[BGS24] Huck Bennett, Surendra Ghentiyala, and Noah Stephens-Davidowitz. The more the merrier! on the complexity of finding multicollisions, with connections to codes and lattices. *Electron. Colloquium Comput. Complex.*, pages TR24–018, 2024. URL: https://eccc.weizmann.ac.il/report/2024/018.

[BJ12] Samuel R. Buss and Alan S. Johnson. Propositional proofs and reductions between NP search problems. *Ann. Pure Appl. Logic*, 163(9):1163–1182, 2012. doi:10.1016/j.apal.2012.01.015.

[BK94] Samuel R. Buss and Jan Krajíček. An application of Boolean complexity to separation problems in bounded arithmetic. *Proc. London Math. Soc. (3)*, 69(1):1–21, 1994. doi:10.1112/plms/s3-69.1.1.

[BKP18] Nir Bitansky, Yael Tauman Kalai, and Omer Paneth. Multi-collision resistance: a paradigm for keyless hash functions. In *STOC'18—Proceedings of the 50th Annual ACM*

SIGACT Symposium on Theory of Computing, pages 671–684. ACM, New York, 2018. doi:10.1145/3188745.3188870.

[BM04] Josh Buresh-Oppenheim and Tsuyoshi Morioka. Relativized NP search problems and propositional proof systems. In *19th Annual IEEE Conference on Computational Complexity (CCC 2004), 21-24 June 2004, Amherst, MA, USA*, pages 54–67. IEEE Computer Society, 2004. doi:10.1109/CCC.2004.1313795.

[CFGH19] David Conlon, Jacob Fox, Andrey Grinshpun, and Xiaoyu He. Online Ramsey numbers and the subgraph query problem. In *Building bridges II—mathematics of László Lovász*, volume 28 of *Bolyai Soc. Math. Stud.*, pages 159–194. Springer, Berlin, [2019] ©2019.

[CG88] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. volume 17, pages 230–261. 1988. Special issue on cryptography. doi:10.1137/0217015.

[CGMS23] Marcelo Campos, Simon Griffiths, Robert Morris, and Julian Sahasrabudhe. An exponential improvement for diagonal Ramsey, 2023. arXiv:2303.09521.

[Coh16] Gil Cohen. Two-source dispersers for polylogarithmic entropy and improved Ramsey graphs. In *STOC'16—Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 278–284. ACM, New York, 2016. doi:10.1145/2897518.2897530.

[CZ19] Eshan Chattopadhyay and David Zuckerman. Explicit two-source extractors and resilient functions. *Ann. of Math. (2)*, 189(3):653–705, 2019. doi:10.4007/annals.2019.189.3.1.

[Das19] Constantinos Daskalakis. Equilibria, fixed points, and computational complexity. In *Proceedings of the International Congress of Mathematicians (ICM). World Scientific*, 2019.

[DP11] Constantinos Daskalakis and Christos H. Papadimitriou. Continuous local search. In Dana Randall, editor, *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 790–804. SIAM, 2011. doi:10.1137/1.9781611973082.62.

[Erd47] Paul Erdös. Some remarks on the theory of graphs. *Bulletin of the American Mathematical Society*, 53:292–294, 1947. URL: https://api.semanticscholar.org/CorpusID:14215209.

[FGHS23] John Fearnley, Paul Goldberg, Alexandros Hollender, and Rahul Savani. The complexity of gradient descent: CLS = PPAD ∩ PLS. *J. ACM*, 70(1):7:1–7:74, 2023. doi:10.1145/3568163.

[FGPR24] Noah Fleming, Stefan Grosser, Toniann Pitassi, and Robert Robere. Black-Box PPP Is Not Turing-Closed. In *STOC'24—Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1405–1414. ACM, New York, 2024. doi:10.1145/3618260.3649769.

[FKP19] Noah Fleming, Pravesh Kothari, and Toniann Pitassi. Semialgebraic proofs and efficient algorithm design. *Found. Trends Theor. Comput. Sci.*, 14(1-2):front matter, 1–221, 2019. doi:10.1561/0400000086.

[GHJ+22] Mika Göös, Alexandros Hollender, Siddhartha Jain, Gilbert Maystre, William Pires, Robert Robere, and Ran Tao. Separations in Proof Complexity and TFNP. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 1150–1161. IEEE, 2022. doi:10.1109/FOCS54457.2022.00111.

[GKRS19] Mika Göös, Pritish Kamath, Robert Robere, and Dmitry Sokolov. Adventures in monotone complexity and TFNP. In *10th Innovations in Theoretical Computer Science*, volume 124 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 38, 19. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019.

[GM08] Konstantinos Georgiou and Avner Magen. Limitations of the Sherali-Adams Lift and Project System: Compromising Local and Global Arguments. Technical Report 587, University of Toronto, 2008.

[GP17] Paul W. Goldberg and Christos H. Papadimitriou. TFNP: an update. In Dimitris Fotakis, Aris Pagourtzis, and Vangelis Th. Paschos, editors, *Algorithms and Complexity - 10th International Conference, CIAC 2017*, volume 10236 of *Lecture Notes in Computer Science*, pages 3–9, 2017. doi:10.1007/978-3-319-57586-5_1.

[GRS23] Venkatesan Guruswami, Atri Rudra, and Madhu Sudan. *Essential Coding Theory (Draft)*. 2023. URL: https://cse.buffalo.edu/faculty/atri/courses/coding-theory/book/.

[HHRS15] Iftach Haitner, Jonathan J. Hoch, Omer Reingold, and Gil Segev. Finding collisions in interactive protocols - tight lower bounds on the round and communication complexities of statistically hiding commitments. *SIAM J. Comput.*, 44(1):193–242, 2015. doi:10.1137/130938438.

[HKT24] Pavel Hubáček, Erfan Khaniki, and Neil Thapen. TFNP intersections through the lens of feasible disjunction. In *15th Innovations in Theoretical Computer Science Conference*, volume 287 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 63, 24. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2024. doi:10.4230/lipics.itcs.2024.63.

[Hou15] Beach House. *PPP*. Aug 2015. URL: https://en.wikipedia.org/wiki/PPP_(song).

[Jeř09] Emil Jeřábek. Approximate counting by hashing in bounded arithmetic. *J. Symbolic Logic*, 74(3):829–860, 2009. doi:10.2178/jsl/1245158087.

[Jeř16] Emil Jeřábek. Integer factoring and modular square roots. *J. Comput. System Sci.*, 82(2):380–394, 2016. doi:10.1016/j.jcss.2015.08.001.

[Jou04] Antoine Joux. Multicollisions in iterated hash functions. application to cascaded constructions. In Matthew K. Franklin, editor, *Advances in Cryptology - CRYPTO 2004, 24th Annual International CryptologyConference, Santa Barbara, California, USA, August 15-19, 2004, Proceedings*, volume 3152 of *Lecture Notes in Computer Science*, pages 306–316. Springer, 2004. doi:10.1007/978-3-540-28628-8_19.

[JPY88] David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? volume 37, pages 79–100. 1988. 26th IEEE Conference on Foundations of Computer Science (Portland, OR, 1985). doi:10.1016/0022-0000(88)90046-3.

[KNY18] Ilan Komargodski, Moni Naor, and Eylon Yogev. Collision resistant hashing for paranoids: dealing with multiple collisions. In *Advances in cryptology—EUROCRYPT 2018. Part II*, volume 10821 of *Lecture Notes in Comput. Sci.*, pages 162–194. Springer, Cham, 2018. URL: https://doi.org/10.1007/978-3-319-78375-8_6, doi:10.1007/978-3-319-78375-8_6.

[KNY19] Ilan Komargodski, Moni Naor, and Eylon Yogev. White-box vs. black-box complexity of search problems: Ramsey and graph property testing. *J. ACM*, 66(5), jul 2019. doi:10.1145/3341106.

[KoNT11] Leszek Aleksander Kołodziejczyk, Phuong Nguyen, and Neil Thapen. The provably total NP search problems of weak second order bounded arithmetic. *Ann. Pure Appl. Logic*, 162(6):419–446, 2011. doi:10.1016/j.apal.2010.12.002.

[KoT22] Leszek Aleksander Kołodziejczyk and Neil Thapen. Approximate counting and NP search problems. *J. Math. Log.*, 22(3):Paper No. 2250012, 31, 2022. doi:10.1142/S021906132250012X.

[Kra05] Jan Krajíček. Structured pigeonhole principle, search problems and hard tautologies. *Journal of Symbolic Logic*, 70(2):619 – 630, 2005. doi:10.2178/jsl/1120224731.

[Li23] Xin Li. Two source extractors for asymptotically optimal entropy, and (many) more. *CoRR*, abs/2303.06802, 2023. arXiv:2303.06802, doi:10.48550/arXiv.2303.06802.

[Li24] Jiawei Li. Total NP search problems with abundant solutions. In *15th Innovations in Theoretical Computer Science Conference*, volume 287 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 75, 23. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2024. doi:10.4230/lipics.itcs.2024.75.

[LZ19] Qipeng Liu and Mark Zhandry. On finding quantum multi-collisions. In *Advances in cryptology—EUROCRYPT 2019. Part III*, volume 11478 of *Lecture Notes in Comput. Sci.*, pages 189–218. Springer, Cham, 2019. doi:10.1007/978-3-030-17659-4.

[Mor01] Tsuyoshi Morioka. Classification of search problems and their definability in bounded arithmetic. *Electron. Colloquium Comput. Complex.*, TR01-082, 2001. URL: https://eccc.weizmann.ac.il/eccc-reports/2001/TR01-082/index.html, arXiv:TR01-082.

[MPW00] Alexis Maciel, Toniann Pitassi, and Alan R. Woods. A new proof of the weak pigeonhole principle. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, pages 368–377. ACM, New York, 2000. doi:10.1145/335305.335348.

[Mül21] Moritz Müller. Typical forcings, NP search problems and an extension of a theorem of riis. *Ann. Pure Appl. Log.*, 172(4):102930, 2021. URL: https://doi.org/10.1016/j.apal.2020.102930, doi:10.1016/J.APAL.2020.102930.

[Pap94] Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. volume 48, pages 498–532. 1994. 31st Annual Symposium on Foundations of Computer Science (FOCS) (St. Louis, MO, 1990). doi:10.1016/S0022-0000(05)80063-7.

[PPY23] Amol Pasarkar, Christos H. Papadimitriou, and Mihalis Yannakakis. Extremal combinatorics, Iterated Pigeonhole Arguments and Generalizations of PPP. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference, ITCS*

2023, volume 251 of *LIPIcs*, pages 88:1–88:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.ITCS.2023.88.

[Ram29] F. P. Ramsey. On a Problem of Formal Logic. *Proc. London Math. Soc. (2)*, 30(4):264–286, 1929. doi:10.1112/plms/s2-30.1.264.

[RV22] Ron D. Rothblum and Prashant Nalini Vasudevan. Collision-resistance from multi-collision-resistance. In *Advances in cryptology—CRYPTO 2022. Part III*, volume 13509 of *Lecture Notes in Comput. Sci.*, pages 503–529. Springer, Cham, [2022] ©2022. URL: https://doi.org/10.1007/978-3-031-15982-4_17, doi:10.1007/978-3-031-15982-4_17.

[ST11] Alan Skelley and Neil Thapen. The provably total search problems of bounded arithmetic. *Proc. Lond. Math. Soc. (3)*, 103(1):106–138, 2011. doi:10.1112/plms/pdq044.

[Tha02] Neil Thapen. A model-theoretic characterization of the weak pigeonhole principle. *Ann. Pure Appl. Logic*, 118(1-2):175–195, 2002. doi:10.1016/S0168-0072(02)00038-6.

[Vad12] Salil P. Vadhan. Pseudorandomness. *Found. Trends Theor. Comput. Sci.*, 7(1-3):1–336, 2012. doi:10.1561/0400000010.

[YZ22] Takashi Yamakawa and Mark Zhandry. Verifiable quantum advantage without structure. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 69–74. IEEE, 2022. doi:10.1109/FOCS54457.2022.00014.

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il