

Multi-pseudodeterministic algorithms*

Oded Goldreich[†]

January 31, 2025

To a great woman.

Abstract

In this work, dedicated to Shafi Goldwasser, we consider a relaxation of the notion of pseudodeterministic algorithms, which was put forward by Gat and Goldwasser (*ECCC*, TR11-136, 2011).

Pseudodeterministic algorithms are randomized algorithms that solve search problems by almost always providing the same canonical solution (per each input). Multi-pseudodeterministic algorithms relax the former notion by allowing the algorithms to output one of a bounded number of canonical solutions (per each input). We show that efficient multi-pseudodeterministic algorithms can solve natural problems that are not solvable by efficient pseudodeterministic algorithms, present a composition theorem regarding multi-pseudodeterministic algorithms, and relate them to other known notions.

A preliminary version of this work appeared as TR19-012 of *ECCC*. Unfortunately, the original presentation contained several inaccuracies and minor errors, which are corrected in the current version. In general, the current revision is more detailed. Although quite a few studies of pseudodeterministic algorithms have appeared since then, we refrain from surveying most of them.

1 Introduction

Randomized algorithms have a clear deficiency: They require a source of randomness (or pseudorandomness), and carry a probability of error (or failure). However, as noted by Gat and Goldwasser [2], in the context of solving search problems, randomized algorithms have another deficiency: They do not necessarily return the same answer, when invoked on the same input. At the extreme, which does occur in many natural algorithms, their output is a random variable that has very low (e.g., exponentially small) collision probability.

Randomized algorithms for solving search problems that avoid the latter deficiency (but rather have very high collision probability) are thus of natural interest. These algorithms offer a functionality that is closer in spirit to that of deterministic algorithms, and are thus called *pseudodeterministic*. That is, pseudodeterministic algorithms are randomized algorithms that solve search problems by almost always providing the same canonical solution.

*Partially supported by the Israel Science Foundation (grant No. 1146/18).

[†]Faculty of Mathematics and Computer Science, Weizmann Institute of Science, Rehovot, ISRAEL. Email: oded.goldreich@weizmann.ac.il.

Definition 1.1 (pseudodeterministic algorithms [2]): For a binary relation $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$, let $R(x) \stackrel{\text{def}}{=} \{y : (x, y) \in R\}$ and $S_R = \{x : R(x) \neq \emptyset\}$. A randomized algorithm A is said to be a pseudodeterministic solver of the search problem R if for every $x \in S_R$ there exists (a canonical solution) $c_x \in R(x)$ such that $\Pr[A(x) = c_x] \geq 2/3$ (and $\Pr[A(x) = \perp] \geq 2/3$ for every $x \notin S_R$).¹

Note that error reduction is applicable to Definition 1.1; that is, for every $\eta: \mathbb{N} \rightarrow (0, 1/3)$, the constant $2/3$ can be replaced by $1 - \eta(|x|)$ if we are willing to repeat the algorithm for $O(\log(1/\eta(|x|)))$ times. (In contrast, error reduction is not applicable to general randomized algorithms for search problems.)

The notion of pseudodeterministic algorithms was put forward by Gat and Goldwasser [2], who initiated their study. In particular, they presented polynomial-time pseudodeterministic algorithms for several natural search problems, which were previously known to have probabilistic polynomial-time solvers, and characterized the class of search problems having polynomial-time pseudodeterministic algorithms (see Theorem 4.2). The study of pseudodeterministic algorithms was extended to the sublinear-time model in [6], and to randomized-NC in [8].

We note that it may be that randomization does not help in the context of solving search problem in polynomial-time. This is certainly the case if the promise problem version of \mathcal{BPP} equals the promise version of \mathcal{P} (see Theorem 4.3). But as long as the latter equality is not known or assumed, pseudodeterministic algorithms are positioned between deterministic and general probabilistic algorithms. In that case, moving from general probabilistic algorithms to pseudodeterministic ones is a step forward, which may even be viewed as a step towards full derandomization. The same holds with respect to the deterministic, pseudodeterministic, and (general) probabilistic versions of \mathcal{NC} . In contrast, in the context of sublinear algorithms, typically, probabilistic algorithms are significantly stronger than deterministic ones. Unfortunately, the results of [6] indicate that typically pseudodeterministic algorithms are not significantly stronger than deterministic ones.

In light of the foregoing, augmenting the collection of known pseudodeterministic algorithms is of major importance. Having failed to do so, led us to an attempt to relax the notion of pseudodeterministic algorithms. Indeed, the contents of this work is the presentation of a relaxation of the notion of pseudodeterministic algorithms, and a study of some of its features. In particular, we observe that in the context of sublinear algorithms this relaxation “buys” power (see Proposition 2.4). Loosely speaking, the relaxation consists of allowing the algorithm to output one of a bounded number of canonical solutions rather than outputting a single canonical solution. Hence, we call such algorithms multi-pseudodeterministic (or m -pseudodeterministic, where $m: \mathbb{N} \rightarrow \mathbb{N}$ is the aforementioned bound).

Contents and Organization. Section 2 evolves around notions: The basic new notion appears in Section 2.1, and related notions are discussed in Sections 2.2 and 2.3. In Section 2.1 we also demonstrate the benefit of multi-pseudodeterministic algorithms in the context of sublinear-time algorithms. In Section 3 we present a composition result for multi-pseudodeterministic algorithms. Loosely speaking, it asserts that *one can solve t instances of a problem that has an efficient m -pseudodeterministic algorithm by using an efficient $(t \cdot (m-1) + 1)$ -pseudodeterministic algorithm*. In Section 4 we characterize the class of search problems that can be solved by polynomial-time multi-pseudodeterministic algorithms in terms of restricted (deterministic polynomial-time) reductions to the promise problem version of \mathcal{BPP} . Loosely speaking, the restriction corresponds to the number

¹The special symbol $\perp \notin \{0, 1\}^*$ indicates failure to find a solution.

of possible leaves in the tree of all possible executions of the deterministic reduction (on a fixed input). In Section 5, we tell the story behind this paper.

The computational models and complexity measures. Our main results, presented in Sections 2 and 3, are applicable to several models of computation and complexity classes such as the standard models of (probabilistic) polynomial-time algorithms and sublinear-time algorithms. This is the case because our results are proved by transformations from one type of algorithm to another type, where the latter algorithm invokes the former algorithm as a subroutine (several times). These transformations presume a model of computation in which the complexity (or cost) of these invocations dominates the complexity of the resulting algorithms. This holds in the aforementioned standard models, but not in the context of log-space computation (e.g., since storing the outcomes of numerous subroutine calls is not for free).

2 Notions

Shafi: *Do you have the notion of a refill?*
The attendant: *Yes, we have refills, but what is a notion?*
(In a diner, on the way to STOC'85.)

We present three alternative quantitative relaxations of the notion of pseudodeterministic algorithms. The first two alternatives are equivalent up to an overhead (in time complexity) that is polynomial in the parameter that governs the relaxation, whereas the equivalence to the third alternative is up to a multiplicative factor in this parameter itself. In Sections 3 and 4 we shall only use the initial definition discussed in Section 2.1. The alternative definitions are presented in Sections 2.2 and 2.3.

2.1 The Basic Notion

One basic difficulty that frustrates attempts to construct pseudodeterministic algorithms is their failure to solve the problem of estimating the average value of a (bounded) function defined over a huge universe, which is easily solvable by ordinary probabilistic algorithms. The issue is that we can easily find an approximate value, and (w.h.p.) this value can be made to reside in a small interval, but we cannot hit a single (canonical) value (with high probability). On the other hand, if we were allowed to hit one of two canonical values (with high probability), then we can easily do it (see Proposition 2.4). This leads to the following generalization of Definition 1.1, where the canonical solution c_x is replaced by a set of at most $m(|x|)$ solutions, denoted C_x . (However, the algorithm is required to hit this set with probability at least $\frac{m(|x|)+1}{m(|x|)+2}$ (rather than with probability at least $\frac{2}{3}$); this choice allows to distinguish between some valid solution and all invalid solutions (as further discussed below).)

Definition 2.1 (*m*-pseudodeterministic algorithm): *For R and S_R as in Definition 1.1, and $m : \mathbb{N} \rightarrow \mathbb{N}$, we say that A is a *m*-pseudodeterministic solver of the search problem R if for every $x \in S_R$ there exists a non-empty set $C_x \subseteq R(x)$ such that $|C_x| \leq m(|x|)$ and $\Pr[A(x) \in C_x] \geq \frac{m(|x|)+1}{m(|x|)+2}$. Furthermore, as in Definition 1.1, it is required that $\Pr[A(x) = \perp] \geq 2/3$ for every $x \notin S_R$.*

(Definition 1.1 is recovered by setting $m \equiv 1$.) As shown in Remark 2.3, error reduction is possible here too, since we can distinguish an output not in C_x (which occurs with probability at most $\frac{1}{m+2}$) from at least one of the outputs in C_x (which occurs with probability at least $\frac{1+(1/m)}{m+2}$).² Indeed, as reflected in the following result, the gap between the probability of outputting some element in $R(x)$ and the probability of outputting any element outside $R(x)$ is a salient feature of Definition 2.1.

Proposition 2.2 (key feature of m -pseudodeterministic algorithms): *Let R, S_R and $m : \mathbb{N} \rightarrow \mathbb{N}$ be as in Definition 2.1. Then, any m -pseudodeterministic solver of R , denoted A , satisfies*

1. *For every $x \in S_R$, there exists $c_x \in R(x)$ such that $\Pr[A(x) = c_x] \geq \frac{1+(1/m)}{m+2}$ and for every $y \notin R(x)$ it holds that $\Pr[A(x) = y] \leq 1/(m+2)$.*
2. *For every $x \notin S_R$, it holds that $\Pr[A(x) = \perp] \geq 2/3$.*

On the other hand, if an algorithm A satisfies the foregoing two conditions (w.r.t R), then there exists an m -pseudodeterministic solver of R that invokes A for $\tilde{O}(m^6)$ times.

Condition 2 appears explicitly in both Definition 1.1 and Definition 2.1. Hence, the difference between the two definitions is actually captured by the probability bounds that appear in Condition 1 of Proposition 2.2: An m -pseudodeterministic algorithm outputs some valid solution with probability that is noticeably larger than $1/(m+1)$, while outputting every invalid solution with probability that is noticeably smaller than $1/(m+1)$, where the threshold of $1/(m+2)$ is chosen for simplicity (see Footnote 2).

Proof: The main claim follows by an averaging argument; that is, for any $x \in S_R$, it holds that $\max_{y \in C_x} \{\Pr[A(x) = y]\} \geq \Pr[A(x) \in C_x]/|C_x|$, where C_x is as in Definition 2.1.

Turning to the reverse claim (i.e., the opposite direction), suppose that A satisfies the two conditions, and consider an algorithm A' that invokes A for $t \stackrel{\text{def}}{=} \tilde{O}(m^6)$ times and outputs the most frequently occurring outcome. To show that algorithm A' is m -pseudodeterministic, we let Y_x be a random variable that represents the set of strings that occur in at least $\tau \stackrel{\text{def}}{=} \frac{1+(1/m)-(1/3m^2)}{m+2} \cdot t$ of the t invocations of $A(x)$ (in a random execution of $A'(x)$); that is, letting $A(x, r)$ denote the output of A , on input x and randomness r , it holds that $y \in Y_x(r_1, \dots, r_t)$ if and only if $|\{i \in [t] : A(x, r_i) = y\}| \geq \tau$. We now make the following four observations.

(i) If $\Pr[A(x) = y] \geq \frac{1+(1/m)}{m+2}$, then $\Pr[y \in Y_x] = 1 - o(1/m)$.

(This is the case because $\frac{\tau}{t} = \frac{1+(1/m)}{m+2} - \frac{1/3m^2}{m+2}$.)

(Indeed, here and in (ii), we use $t \geq O(m^6 \log m)$.)

(ii) Let $H_x \stackrel{\text{def}}{=} \left\{ y : \Pr[A(x) = y] \geq \frac{\tau}{t} - \frac{1}{6(m+2)^3} \right\}$.

Then, $\Pr[Y_x \subseteq H_x] = 1 - o(1/m)$.

(Recalling that $\tau = \Theta(t/m)$, partition the elements that are not in H_x into $O(m)$ sets, denoted $S_1, \dots, S_{O(m)}$, such that $\Pr[A(x) \in S_i] < (\tau/t) - (1/6(m+2)^3)$ for each i , and observe that $\Pr[S_i \cap Y_x \neq \emptyset] = 1 - o(1/m^2)$.)

²We use $\frac{m+1}{m+2}$ (resp., $\frac{1}{m+2}$) as a quantity that is noticeably larger than $\frac{m}{m+1}$ (resp., noticeably smaller than $\frac{1}{m+1}$). In general, the complexity of error-reduction is polynomially related to the reciprocal of this noticeable gap.

(iii) By the hypothesis, $c_x \in H_x$ for any $x \in S_R$.

(iv) As shown next, $|H_x| \leq m$ for any x .

This is proved by observing that otherwise $\Pr[A(x) \in H_x] > 1$. Specifically, the contradiction hypothesis implies that $\sum_{y \in H_x} \Pr[A(x) = y]$ is at least

$$\begin{aligned} (m+1) \cdot \left(\frac{\tau}{t} - \frac{1}{6(m+2)^3} \right) &= (m+1) \cdot \left(\frac{1 + (1/m) - (1/3m^2)}{m+2} - \frac{1}{6(m+2)^3} \right) \\ &= 1 - \frac{(2m-1)/3m}{m(m+2)} - \frac{m+1}{6(m+2)^3} > 1. \end{aligned}$$

Combining Observations (iii) and (i), $\Pr[c_x \in Y_x] = 1 - o(1/m)$ holds for any $x \in S_R$. Hence, in that case, $Y_x \neq \emptyset$ and so $A'(x)$ is in Y_x . On the other hand, by Observation (ii), $\Pr[Y_x \subseteq H_x] = 1 - o(1/m)$, whereas $|H_x| \leq m$ by Observation (iv). Hence, in that case, $A'(x)$ is one of at most m strings. Furthermore, observing that $H_x \subseteq R(x)$, since $\tau > 1/(m+2)$, it follows that $\Pr[A'(x) \in R(x)] = 1 - o(1/m)$ holds for any $x \in S_R$. ■

Remark 2.3 (error reduction for m -pseudodeterministic algorithm): *As is evident from the proof of Proposition 2.2, the error probability of m -pseudodeterministic algorithm can be reduced from $1/(m+2)$ to $\eta < 1/(m+2)$ by using $t = O(m^6 \log(1/\eta))$ repetitions of the original algorithm. Under this setting, for each constant $c \geq 1$, the phrase “with probability $1 - o(1/m^c)$ ” used in the proof can be replaced by “with probability $1 - o(\eta/m^{c-1})$ ”.*

A demonstration of the benefit of the generalization.

Indeed, as asserted in the motivating discussion, the generalization from pseudodeterministic algorithms to 2-pseudodeterministic algorithms allows for approximating the average value of a bounded function defined over a huge universe.

Proposition 2.4 (averages can be approximated by a 2-pseudodeterministic algorithm): *Given oracle access to $f : \{0,1\}^n \rightarrow [0,1]$, the value $\bar{f} \stackrel{\text{def}}{=} \mathbb{E}_x[f(x)]$ can be approximated, with high constant probability, to within an additive error of ϵ by a 2-pseudodeterministic algorithm that makes $O(1/\epsilon^2)$ queries to f .*

Recall that (by [6, Thm. 4.1]) any 1-pseudodeterministic algorithm that solves this problem must have query complexity $\exp(\Omega(n))$.

Proof: The algorithm selects uniformly at random $m = O(1/\epsilon^2)$ points $x_1, \dots, x_m \in \{0,1\}^n$, queries f for their value, and outputs $\lfloor v/\epsilon \rfloor \cdot \epsilon$ such that $v = \sum_{i \in [m]} f(x_i)/m$, where $\lfloor \alpha \rfloor$ is the integer closest to α . The claim follows by noting that, with probability at least 0.9, it holds that $|v - \bar{f}| < 0.1\epsilon$, whereas $|\lfloor v/\epsilon \rfloor \cdot \epsilon - v| \leq 0.5\epsilon$ always holds. In this case (i.e., when $|v - \bar{f}| < 0.1\epsilon$), we have $\lfloor v/\epsilon \rfloor \in [\lfloor (\bar{f} - 0.1\epsilon)/\epsilon \rfloor, \lfloor (\bar{f} + 0.1\epsilon)/\epsilon \rfloor] = [\lfloor \bar{f}/\epsilon \rfloor - 0.1, \lfloor \bar{f}/\epsilon \rfloor + 0.1]$, whereas the latter interval contains at most two integers (since its length is at most 1). ■

2.2 An Alternative Definition

The following definition, which is related to Definition 2.1, was suggested by Salil Vadhan. It interprets m -pseudodeterministic algorithms as ones capable of outputting (w.h.p.) a list of m valid solutions that always contains the canonical solution (where c_x denotes the canonical solution for x). In the following definition, $A(x)$ represents a set (or list) of potential solutions, and $|A(x)|$ denotes the size of this set (resp., list).

Definition 2.5 (m -pseudodeterministic algorithm, alternative): *For R, S_R and $m : \mathbb{N} \rightarrow \mathbb{N}$ as in Definition 2.1, we say that A is a m -pseudodeterministic list solver of the search problem R if for every $x \in S_R$ there exists $c_x \in R(x)$ such that $\Pr[c_x \in A(x) \subseteq R(x)] \geq 1 - \frac{1}{3m(|x|)}$ and $|A(x)| \leq m(|x|)$ always holds. Furthermore, as in Definition 2.1, it is required that $\Pr[A(x) = \perp] \geq 2/3$ for every $x \notin S_R$.*

(As in Definition 2.1, Definition 1.1 is recovered by setting $m \equiv 1$.) We stress that $A(x)$ may have very low collision probability (equiv., $\Pr[A(x) = S]$ may be exponentially vanishing (in $|x|$) for every set S). As elaborated in Remark 2.8, the choice of the probability bound in the main condition of Definition 2.5 is more arbitrary than in Definition 2.1.

Theorem 2.6 (relating Definitions 2.1 and 2.5): *For every efficiently computable $m : \mathbb{N} \rightarrow \mathbb{N}$ and search problem R , the following hold.*

1. *If the search problem R has an m -pseudodeterministic solver, then R has an m -pseudodeterministic list solver that invokes the original solver $\tilde{O}(m^4)$ times.*
2. *If the search problem R has an m -pseudodeterministic list solver, then R has an m -pseudodeterministic solver that invokes the original solver $\tilde{O}(m^2)$ times.*

In both cases, the invocations of the original solvers dominate the complexity of the resulting solver.

Proof: Starting with Part 1, let A be an m -pseudodeterministic solver of R . For an arbitrary function $\eta : \mathbb{N} \rightarrow (0, 1/3]$, consider an algorithm A' that on input x proceeds as follows.

1. Invokes $A(x)$ for $t \stackrel{\text{def}}{=} O(m(|x|)^4 \cdot \log(1/\eta(|x|)))$ times, and let p_y denote the frequency of the output y in these t invocations; that is, p_y equals $|\{i \in [t] : y_i = y\}|/t$, where y_i denotes the output of the i^{th} invocation of $A(x)$.
2. If $p_\perp > 1/2$, then we output \perp . Otherwise, we output the list of all y 's such that $p_y > \frac{1+(1/2m(|x|))}{m(|x|)+2}$.

Clearly, if $x \notin S_R$, then $A'(x)$ equals \perp with very high probability. Turning to the case of $x \in S_R$, let C_x be as in Definition 2.1, and let $c_x \in C_x$ be a string that is output by $A(x)$ with the highest probability (with ties broken arbitrarily). Then (as shown in the first part of the proof of Proposition 2.2), $\Pr[A(x) = c_x] \geq \frac{1+(1/m(|x|))}{m(|x|)+2}$, and so, with probability at least $1 - \eta(|x|)/2$ over the executions of $A'(x)$, it holds that $p_{c_x} > \frac{1+(1/2m(|x|))}{m(|x|)+2}$. On the other hand, $\Pr[A(x) \notin C_x] \leq \frac{1}{m(|x|)+2}$, and so, with probability at least $1 - \eta(|x|)/2$, it holds that $\sum_{y \notin C_x} p_y < \frac{1+(1/2m(|x|))}{m(|x|)+2}$. Recalling that $|C_x| \leq m(|x|)$ and $C_x \subseteq R(x)$, it follows that, with probability at least $1 - \eta(|x|)$, the execution

of $A'(x)$ returns a list of size at most $m(|x|)$ that contain c_x and is contained in $R(x)$. Setting $\eta(n) = 1/3m(n)$, Part 1 follows (since we may output \perp in the rare case that the aforementioned list exceeds the size bound).

Turning to Part 2, let A be an m -pseudodeterministic list solver of R , and consider the probability, denoted q_y , that y occurs in the list output by $A(x)$; that is, $q_y \stackrel{\text{def}}{=} \Pr[y \in A(x)]$. Then, $\sum_y q_y = \text{Exp}[|A(x)|] \leq m(|x|)$, which implies $|\{y : q_y > 1 - 1/(m(|x|) + 1)\}| \leq m(|x|)$, since otherwise we reach a contradiction (i.e., $\sum_y q_y > (m + 1) \cdot (1 - 1/(m + 1)) = m$). Furthermore, by the hypothesis, $q_{c_x} \geq 1 - 1/3m(|x|) > 1 - 1/(m(|x|) + 1)$, and $q_y \leq 1/3m(|x|) < 1 - 1/(m(|x|) + 1)$ for every $y \notin R(x)$. Now, consider an algorithm A' that on input x proceeds as follows.

1. Invokes $A(x)$ for $t \stackrel{\text{def}}{=} O(m(|x|)^2 \cdot \log m(|x|))$ times, and let p_y denote the frequency of the occurrence of y in the lists produces in these t invocations; that is, p_y equals $|\{i \in [t] : y \in L_i\}|/t$, where L_i denotes the list output by the i^{th} invocation of $A(x)$.
2. If some y satisfies $p_y > 1 - 1/2m(|x|)$, then we output it (i.e., if several y 's satisfy $p_y > 1 - 1/2m(|x|)$, then we output one of them chosen arbitrarily). Otherwise, we output \perp .

Clearly, if $x \notin S_R$, then $A'(x)$ equals \perp with very high probability. Otherwise (i.e., $x \in S_R$), with probability at least $1 - o(1/m(|x|))$ over the execution of $A'(x)$, it holds that $p_{c_x} \geq 1 - 1/2m(|x|)$, since $q_{c_x} \geq 1 - 1/3m(|x|)$, and in this case A' output a string (i.e., not \perp). Likewise, with probability at least $1 - o(1/m(|x|))$, no y such that $q_y \leq 1 - 1/(m(|x|) + 1)$ satisfies $p_y \geq 1 - 1/2m(|x|)$.³ Hence, with probability at least $1 - o(1/m(|x|))$, the execution $A'(x)$ outputs a string in $\{y : q_y > 1 - 1/(m(|x|) + 1)\} \subseteq R(x)$. Recalling that the former set has cardinality at most $m(|x|)$, Part 2 follows. ■

Remark 2.7 (a single algorithm that interpolates Definitions 2.1 and 2.5): *The proof of Part 2 of Theorem 2.6 yields subsets $L(x) \subseteq U(x)$ of $R(x)$ and an algorithm A'' such that (i) $|U(x)| \leq m(|x|)$, (ii) $L(x) \neq \emptyset$ for $x \in S_R$, and (iii) $\Pr[L(x) \subseteq A''(x) \subseteq U(x)] = 1 - o(1/m(|x|))$.⁴ While A'' constitutes an m -pseudodeterministic list solver, outputting an arbitrary element in $A''(x)$ yields an m -pseudodeterministic solver.*

Remark 2.8 (on the probability bound in the main condition of Definition 2.5). *The choice of the probability bound in the main condition of Definition 2.5 (i.e., $\Pr[c_x \in A(x) \subseteq R(x)] \geq 1 - \frac{1}{3m(|x|)}$) is more arbitrary than in Definition 2.1. For starters, as in the latter case, any lower bound that is noticeably larger than $m/(m+1)$ (and smaller than $1 - \exp(-m)$) will do (see proof of Theorem 2.6).⁵ Furthermore, even an lower bound that is (only) noticeably larger than $1/2$ allows for equivalence with Definition 2.1, albeit at the cost increasing the parameter m . For example, a lower bound of*

³Here, we assume $m > 1$, since otherwise the claim is trivial. As shown in the proof of Observation (ii) in the proof of Proposition 2.2, the claim regarding such y 's is shown by partitioning all these y 's into sets, $S_1, \dots, S_{O(m)}$, such that $\Pr[A(x) \cap S_i \neq \emptyset] \leq 1 - 1/(m + 1)$ for each i , and arguing separately about the frequency of each S_i hitting the t lists generated by $A'(x)$ (i.e., with probability at least $1 - o(1/m^2)$ over the execution of $A'(x)$, elements of S_i occur in less than $(1 - 1/2m) \cdot t$ of the lists).

⁴Indeed, $U(x) = \{y : \Pr[y \in A(x)] > 1 - 1/(m(|x|) + 1)\}$ and $L(x) = \{y : \Pr[y \in A(x)] > 1 - 1/3m(|x|)\}$.

⁵Specifically, note the free parameter η in the proof of the first direction, and observe that in the opposite direction replacing the lower bound $1 - 1/3m$ by p such that $p > m/(m + 1)$ suffices in order to assert that $c_x \in \{y : q_y > 1 - 1/(m + 1)\} \subseteq R(x)$, whereas the threshold used in Step 2 of algorithm A' (i.e., $1 - 1/2m$) should be replaced by $\frac{p + (m/(m+1))}{2}$. Hence, the choice of p only affects the overhead factor, which is $O((p - (m/(m + 1)))^{-2})$.

2/3 in the main condition of Definition 2.5 allows for an equivalence up to a constant factor in the parameter m ; see the proof of Theorem 2.6.⁶

2.3 Relation to Reproducible Solution Algorithms

The following (general) notion of reproducible solution algorithms is implicit in the work of Grossman and Liu [9], which focuses on the case of randomized log-space. Loosely speaking, an ℓ -reproducible solver consists of an algorithm that (w.h.p.) outputs an ℓ -bit long concise representation of a solution that can be later reproduced in full. Hence, such an algorithm yields a 2^ℓ -pseudodeterministic solver. The following definition mandates a less high probability bound, and so error-reduction is required for the foregoing implication.

Definition 2.9 (ℓ -reproducible algorithms): *For R, S_R as in Definition 2.1 and $\ell : \mathbb{N} \rightarrow \mathbb{N}$, we say that A is an ℓ -reproducible solver of the search problem R if A decomposes to two algorithms, A_1 and A_2 , such that $(A(x) = A_2(x, A_1(x)))$ and the following conditions hold.*

1. $A_1(x) \in \{0, 1\}^{\ell(|x|)}$ always holds, and for $x \notin S_R$ it holds that $\Pr[A_2(x, A_1(x)) = \perp] \geq 2/3$.
2. For every $x \in S_R$ and every r , there exists $c_{x,r} \in R(x)$ such that

$$\text{Exp}_{r \leftarrow A_1(x)} [\Pr[A_2(x, r) = c_{x,r}]] \geq 0.99.$$

Algorithm A_2 (and consequently algorithm A) is **sound** if for every $x \in S_R$ and every r , it holds that $\Pr[A_2(x, r) \in R(x) \cup \{\perp\}] = 1$. Algorithm A is **almost-sound** if, for every $x \in S_R$, it holds that $\Pr[A_2(x, A_1(x)) \in R(x) \cup \{\perp\}] > 1 - 2^{-3\ell(|x|)-7}$.

We say that r is **x -good** if $\Pr[A_2(x, r) = c_{x,r}] \geq 0.9$. Indeed, Condition 2 does not imply that all r 's are good, but rather implies that $A_1(x)$ is x -good with probability at least 0.9. Intuitively, soundness means that A never outputs an invalid solution; a sound algorithm always outputs either a valid solution or a failure symbol (i.e., \perp).

Observe that, like in Definition 2.1, the error probability of algorithm A_2 on inputs (x, r) such that r is x -good can be reduced by repetitions. When A_2 is sound, the error probability of algorithm A_1 can also be reduced by repetitions (and testing the potential outputs by invoking A_2).⁷ A strengthening of the latter fact is used when proving Part 1 of the following result.

Theorem 2.10 (relating Definitions 2.1 and 2.9): *For every efficiently computable $\ell : \mathbb{N} \rightarrow \mathbb{N}$ and search problem R , the following hold.*

1. If R has an ℓ -reproducible almost-sound solver, then, invoking this solver for $O(\ell^2)$ times yields a 2^ℓ -pseudodeterministic solver.

⁶Specifically, in such a case, we shall use the fact that $q_{c_x} \geq 2/3$, whereas $|\{y : q_y > 1/2\}| < 2m$ and $q_y < 1/2$ for every $y \notin R(x)$. Outputting an arbitrary element that appears in at least 60% of the lists, we derive an $2m$ -pseudodeterministic solver. In general, replacing the lower bound $2/3$ by any p such that $p > 1/2$ implies $|\{y : q_y > 1/2\}| < 2m$, whereas the choice of $p > 1/2$ only affects the overhead factor, which is $O((p - 0.5)^{-2})$.

⁷Specifically, we sample a few outputs r 's of $A_1(x)$, and output one that seems to be x -good; that is, we output an r that maximizes $\max_{y \in \{0,1\}^*} \{\Pr[A_2(x, r) = y]\}$, where $\max_{y \in \{0,1\}^*} \{\Pr[A_2(x, r) = y]\} = 0$ is possible for non-good r 's (since we may have $A_2(x, r) \equiv \perp$). We stress that Definition 2.9 does not imply that the maximizing r is necessarily x -good; it may even be that $A_2(x, r) \equiv y$ for some $y \neq c_{x,r}$. Yet, the soundness condition (for $x \in S_R$) implies that the value of $A_2(x, r)$ is always in $R(x) \cup \{\perp\}$. Hence, it is most likely that (for the chosen r) the most frequent value of $A_2(x, r)$ is in $R(x)$, and we may re-define $c_{x,r}$ accordingly.

2. If R has a 2^ℓ -pseudodeterministic solver, then, invoking this solver for $\exp(O(\ell))$ times yields a $(\ell + 2)$ -reproducible almost-sound solver.

Part 2 was suggested to us by Ofer Grossman.

Proof: Starting with Part 1 and employing error reduction, we obtain algorithms A'_1 and A'_2 such that for every $x \in S_R$

$$\Pr_{r \leftarrow A'_1(x)} \left[\Pr[A'_2(x, r) = c_{x,r}] > 1 - 2^{-2\ell(|x|)-1} \right] \geq 1 - 2^{-2\ell(|x|)-1}, \quad (1)$$

where $c_{x,r} \in R(x)$ is as in Definition 2.9. Specifically, $A'_1(x)$ invokes $A_1(x)$ for $t = \ell(|x|) + 3$ times, obtaining the outcomes r_1, \dots, r_t , and outputs $r = r_i$ that maximizes $p_{x,r} \stackrel{\text{def}}{=} \max_{y \in \{0,1\}^*} \{\Pr[A_2(x, r) = y]\}$, where the $p_{x,r}$'s are estimated by invoking $A_2(x, r)$ for $O(\ell(|x|))$ times. In particular, with probability at least $1 - 2^{-2\ell(|x|)-2}$, all p_{x,r_i} 's are estimated up to additive deviation of 0.1. Assuming that $x \in S_R$, note that, with probability at least $1 - 0.1^{\ell(|x|)} > 1 - 2^{-2\ell(|x|)-3}$, at least one of the r_i 's is x -good, and in this case $p_{x,r} \geq 0.7$ for the r that A'_1 outputs.⁸ Furthermore, with probability at least $1 - 2^{-2\ell(|x|)-3}$, the most frequent value of $A_2(x, r)$ is in $R(x)$, since (by the almost-soundness of A) all sampled r_i 's satisfy $\Pr[A_2(x, r) \in R(x) \cup \{\perp\}] > 0.9$ with probability at least $1 - \ell(|x|) \cdot (2^{-3\ell(|x|)-7}/0.1) > 1 - 2^{-2\ell(|x|)-3}$. By possibly re-defining $c_{x,r}$ (for r 's that are not x -good such that $c_{x,r}$ equals the most frequent value of $A_2(x, r)$) as well as relaxing the definition of x -good (so that $c_{x,r}$ is output with probability at least 0.7), we conclude that $A'_1(x)$ outputs an x -good r with probability at least $1 - 2^{-2\ell(|x|)-1}$. Likewise, $A'_2(x, r)$ invokes $A_2(x, r)$ for $O(\ell(|x|))$ times, and outputs the most frequently occurring output. Hence, if r is x -good, then $\Pr[A_2(x, r) = c_{x,r}] > 1 - 2^{-2\ell(|x|)-1}$.

Letting $C_x = \{c_{x,r} : r \in \{0,1\}^{\ell(|x|)}\} \cap R(x)$, we consider an algorithm that on input x outputs $A'(x) \stackrel{\text{def}}{=} A'_2(x, A'_1(x))$. Observe that on input $x \in S_R$, the output of $A_1(x)$ is an x -good with probability at least $1 - 2^{-2\ell(|x|)-1}$, and in that case $A'(x)$ is in C_x with probability at least $1 - 2^{-2\ell(|x|)-1}$. Hence, $A'(x)$ outputs an element of C_x with probability at least $(1 - 2^{-2\ell(|x|)-1})^2 > 1 - 1/(2^{\ell(|x|)} + 2)$. Part 1 follows (by noting that $|C_x| \leq 2^{\ell(|x|)}$ and that the algorithm outputs \perp w.h.p whenever $x \notin S_R$).

Turning to Part 2, we are given an algorithm A that on input $x \in S_R$ satisfies $\Pr[A(x) \in C_x] \geq \frac{m+1}{m+2}$, where C_x is as in Definition 2.1 and $m = 2^{\ell(|x|)}$. The basic idea is distinguishing between elements in C_x that each occur in $A(x)$ with probability at least $\frac{1+(1/2m)}{m+2}$ and the other elements.⁹ The problem is that some elements of C_x may occur with probability that is very close to this threshold. This problem is resolved by using a random threshold τ in the interval $\left[\frac{1+(1/2m)}{m+2} \pm \frac{1/3m}{m+2}\right]$. Specifically, we select the threshold τ uniformly at random in the set $T = \left\{ \frac{1+(1/6m)}{m+2} + \frac{(i-0.5)/6m^2}{m+2} : i \in [4m] \right\}$.

Using the fact that the elements in T are at distance at least $\delta = \frac{1}{6m^2 \cdot (m+2)}$ apart, it follows that each element in $P = \{\Pr[A(x) = y] : y \in C_x\}$ can be at distance at most $\delta/3$ from at most one of the elements of T . Hence, at most m elements of P are $\delta/3$ -close to T , and so with probability

⁸In this case, for r_i that is x -good, it holds that $p_{x,r_i} \geq 0.9$, whereas the corresponding estimate value is at least 0.8. On the other hand, for the maximizing r it holds that $p_{x,r} \geq 0.7$, because its estimate value is at least 0.8.

⁹Recall that (as shown in Proposition 2.2) some string in C_x appears in $A(x)$ with probability at least $\frac{1+(1/m)}{m+2}$, whereas each string outside C_x appears in $A(x)$ with probability at most $\frac{1}{m+2}$.

at least $3/4$, the random threshold $\tau \in T$ is at distance greater than $\frac{1}{18m^2 \cdot (m+2)}$ from each element of P . Note that τ is always at distance at least $1/(6m \cdot (m+2))$ from the probability of occurrence of any element outside of C_x , since these elements occur with probability at most $1/(m+2)$. Thus, we obtain the following $(\ell(|x|) + 2)$ -reproducible solver, which consists of the algorithms A'_1 and A'_2 .

- On input x , the first component (i.e., A'_1) outputs a uniformly distributed $i \in [4 \cdot m]$, where $m = 2^{\ell(|x|)}$.
- On input (x, i) , the second component (i.e., A'_2) to determine the threshold $\tau = \frac{1+(1/6m)}{m+2} + \frac{i-0.5}{6 \cdot m^2 \cdot (m+2)}$ and the set $H_{x,\tau} = \{y : \Pr[A(x)=y] > \tau\}$. It then outputs the lexicographically first element of $H_{x,\tau}$ (and \perp if $H_{x,i} = \emptyset$).¹⁰

Whenever the threshold is good (i.e., is at distance greater than $\frac{1}{18m^2 \cdot (m+2)} = \Omega(2^{-3\ell(|x|)})$ from each element in $\{\Pr[A(x)=y] : y \in \{0,1\}^*\}$), the set $H_{x,\tau}$ is correctly reconstructed (from i) with probability $1 - o(1)$, by using $m' \stackrel{\text{def}}{=} O(2^{6\ell(|x|)} \cdot \ell(|x|))$ invocations of $A(x)$. Hence, letting $c_{x,i}$ denote the lexicographically first element of $H_{x,\tau}$, it holds that

$$\Pr_{i \leftarrow A'_1(x)} [\Pr[A'_2(x, i) = c_{x,i}] = 1 - o(1)] \geq 3/4. \quad (2)$$

Using error-reduction on A'_1 , we can replace $3/4$ by 0.999 , and so obtain an $(\ell + 2)$ -reproducible solver. Finally, observe that (A'_1, A'_2) is almost-sound, because, for every τ , the set of elements deemed to be in $H_{x,\tau}$ is a subset of $R(x) \cup \{\perp\}$ with probability at least $1 - \exp(-\Omega(m'/2^{6\ell(|x|)})) > 1 - 2^{-3\ell(|x|)-7}$. Thus, Part 2 follows. ■

Remark 2.11 (an alternative version of Part 2 of Theorem 2.10): *While Eq. (2) falls short of establishing an $(\ell+2)$ -reproducible solver, its first component (i.e., A'_1) outputs a uniform distributed $(\ell+2)$ -bit long string. Using a threshold at random in the set $T' = \left\{ \frac{1+(1/192m)}{m+2} + \frac{(i-0.5)/192m^2}{m+2} : i \in [128m] \right\}$ (rather than in T). This yields an $(\ell+7)$ -reproducible solver with a first component that outputs a uniform distributed $(\ell+7)$ -bit long string. We warn that employing error reduction, as at the end of the proof of Theorem 2.10, may ruin this feature.*

3 Multiple Invocations of Multi-Pseudodeterministic Algorithms

A straightforward application of Proposition 2.4 to the problem of approximating t different quantities yields a 2^t -pseudodeterministic algorithm. We can do better.

Algorithm 3.1 (approximating t averages by a $(t+1)$ -pseudodeterministic algorithm): *Given oracle access to $f_1, \dots, f_t : \{0,1\}^n \rightarrow [0,1]$, the algorithm proceeds as follows.*

1. *It selects uniformly at random $m = \tilde{O}(t^4)/\epsilon^2$ points, $x_1, \dots, x_m \in \{0,1\}^n$, queries each f_i for its value at each point, and computes $v_i = \sum_{j \in [m]} f_i(x_j)/m$ for each $i \in [t]$.*
2. *It selects uniformly at random a small shift $\sigma \in S \stackrel{\text{def}}{=} \{(j-1) \cdot \frac{\epsilon}{16t^2} : j \in [8t^2]\}$, and sets $v'_i = v_i + \sigma$ for each $i \in [t]$.*

¹⁰The lex-first element of the set is used as an arbitrary function that maps sets of strings to a single string in the set.

It outputs the t -tuple $(\lfloor v'_1/\epsilon \rfloor \cdot \epsilon, \dots, \lfloor v'_t/\epsilon \rfloor \cdot \epsilon)$.

Outputting $\lfloor v'_i/\epsilon \rfloor \cdot \epsilon$ rather than $\lfloor v_i/\epsilon \rfloor \cdot \epsilon$ means that we use a random threshold when “discretizing” the v_i ’s; that is, rather than using the thresholds $\{(k - 0.5) \cdot \epsilon : k \in [1/\epsilon]\}$ (as in the proof of Proposition 2.4), we effectively use the thresholds $\{(k - 0.5) \cdot \epsilon - \sigma : k \in [1/\epsilon]\}$.

Claim 3.2 (analysis of Algorithm 3.1): *For any $t \geq 1$, Algorithm 3.1 constitute a $(t + 1)$ -pseudodeterministic algorithm for approximating the averages of t bounded functions.*

Proof: By the setting of m , with probability at least $1 - \frac{1}{2(t+3)}$, it holds that $|v_i - \bar{f}_i| < 0.01\epsilon/t^2$ for every $i \in [t]$, where $\bar{f}_i \stackrel{\text{def}}{=} \text{Exp}_x[f_i(x)]$, and in this case for every choice of σ it holds that $\lfloor (v_i + \sigma)/\epsilon \rfloor \in I_i(\sigma)$, where

$$I_i(\sigma) = [\lfloor (\bar{f}_i + \sigma - (0.01\epsilon/t^2))/\epsilon \rfloor, \lfloor (\bar{f}_i + \sigma + (0.01\epsilon/t^2))/\epsilon \rfloor].$$

Note that the interval $I_i(\sigma)$ has length at most 1 (equiv., contains at most two integers), and that it always contains the integer $\lfloor (\bar{f}_i + \sigma)/\epsilon \rfloor$. Hence, if $I_i(\sigma)$ contains a single integer (and $|v_i - \bar{f}_i| < 0.01\epsilon/t^2$), then it holds that $\lfloor (v_i + \sigma)/\epsilon \rfloor = \lfloor (\bar{f}_i + \sigma)/\epsilon \rfloor$.

We call σ **bad** for i if the interval $I_i(\sigma)$ contains more than a single integer (equiv., $\lfloor \epsilon^{-1} \cdot (\bar{f}_i + \sigma) - (0.01/t^2) \rfloor \geq \lfloor \epsilon^{-1} \cdot (\bar{f}_i + \sigma) + (0.01/t^2) \rfloor$). Observe that the probability that σ is bad for some $i \in [t]$ is at most

$$\begin{aligned} & \sum_{i \in [t]} \Pr_{\sigma \in S} \left[\left\lfloor \epsilon^{-1} \cdot (\bar{f}_i + \sigma) - \frac{0.01}{t^2} \right\rfloor \neq \left\lfloor \epsilon^{-1} \cdot (\bar{f}_i + \sigma) + \frac{0.01}{t^2} \right\rfloor \right] \\ & \leq t \cdot \max_{x \in \mathbb{R}} \left\{ \Pr_{j \in [8t^2]} \left[\left\lfloor x + \frac{j}{16t^2} - \frac{0.01}{t^2} \right\rfloor \neq \left\lfloor x + \frac{j}{16t^2} + \frac{0.01}{t^2} \right\rfloor \right] \right\} \end{aligned}$$

Letting $L = x - (0.01/t^2)$ and $H = x + (0.01/t^2)$, note that $\lfloor L + r \rfloor + 1 = i + 1 = \lfloor H + r \rfloor$ if and only if $i + 0.5 \in [L + r, H + r]$. In other words, $\lfloor L + r \rfloor + 1 = \lfloor H + r \rfloor$ if and only if $(\mathbb{Z} + 0.5) \cap [L + r, H + r] \neq \emptyset$. Hence, the foregoing quantity is upper-bounded by

$$t \cdot \max_{y \in \mathbb{R}} \left\{ \Pr_{j \in [8t^2]} \left[(\mathbb{Z} + 0.5) \cap [y + (j/16t^2), y + (0.02/t^2) + (j/16t^2)] \neq \emptyset \right] \right\}$$

which equals $t/8t^2$ because the intervals $[y + (j/16t^2), y + (0.02/t^2) + (j/16t^2)]$ are disjoint and are contained in an interval of length smaller than 1. Hence, with probability $1 - \frac{1}{2(t+3)} - \frac{1}{8t} \geq 1 - \frac{1}{t+3}$, all v_i ’s are $0.01\epsilon/t^2$ -close to the corresponding \bar{f}_i ’s and σ is not bad for any i . In this good case, the output sequence equals $(\lfloor (\bar{f}_1 + \sigma)/\epsilon \rfloor \cdot \epsilon, \dots, \lfloor (\bar{f}_t + \sigma)/\epsilon \rfloor \cdot \epsilon)$.

The key observation is that the foregoing sequence can assume at most $t + 1$ possible values when σ varies in the interval $[0, \epsilon/2]$. To see this, let $\text{rem}_\epsilon(\alpha) = \alpha - \lfloor \alpha/\epsilon \rfloor \cdot \epsilon \in [0, \epsilon)$, and assume (w.l.o.g.) that $\text{rem}_\epsilon(\bar{f}_i) \leq \text{rem}_\epsilon(\bar{f}_{i+1})$ for every $i \in [t - 1]$. In this case, a generic sequence $(\lfloor (\bar{f}_1 + \sigma)/\epsilon \rfloor \cdot \epsilon, \dots, \lfloor (\bar{f}_t + \sigma)/\epsilon \rfloor \cdot \epsilon)$ takes the floor-value on the first $i \in \{0, 1, \dots, t\}$ values and the ceiling-value in the remaining $t - i$ values. This means that there are at most $t + 1$ possibilities, and the claim follows. ■

Comment. Algorithm 3.1 is closely related to the $O(\log(t/\epsilon))$ -reproducible solver used by Grossman and Liu [9] to estimate t different probabilities up to a deviation of ϵ (where these probabilities arise from any randomized algorithm that uses space $O(\log t)$). Using Part 1 of Theorem 2.10, their result yields a $\text{poly}(t/\epsilon)$ -pseudodeterministic algorithm, whereas our analysis yields a $(t+1)$ -pseudodeterministic algorithm. The use of a single threshold here (and in [9]) is reminiscent of the use of a single threshold by Saks and Zhou [10], although the benefit of using a single threshold is different. Specifically, Saks and Zhou [10] use a single threshold in order to economize on randomness, whereas we use it in order to better upper-bound the number of possible rounding-patterns (which would have been 2^t otherwise).

Beyond Algorithm 3.1.

Generalizing the idea that underlies the analysis of Algorithm 3.1, we obtain the following composition result (which improves over the trivial bound of m^t).

Theorem 3.3 (non-adaptive invocations of a m -pseudodeterministic algorithm): *Suppose that A is an m -pseudodeterministic algorithm for solving the search problem R . Then, one can solve any t fixed instances of R by a $(t \cdot (m-1) + 1)$ -pseudodeterministic algorithm that invokes A for $\text{poly}(tm)$ times.*

We stress that the t instances are fixed beforehand (rather than selected adaptively based on the answers to prior instances). Note that this result holds both in the context of sub-linear time algorithms (as in Algorithm 3.1) and in the context of polynomial-time algorithms.

Proof: Let (x_1, \dots, x_t) be a sequence of inputs, and suppose for simplicity that $x_1, \dots, x_t \in S_R$ (since the x_i 's that have no solution are easy to detect).

For every $x \in S_R$, let $p_y(x) = \Pr[A(x) = y]$. Observe that $\max_y \{p_y(x)\} \geq \frac{1}{m} \cdot \frac{m+1}{m+2} = \frac{1}{m+2} + \delta$, where $\delta \stackrel{\text{def}}{=} \frac{1}{(m+2) \cdot m}$, whereas $\{y : p_y(x) > \frac{1}{m+2}\} \subseteq C_x$ (where $C_x \subseteq R(x)$ is as in Definition 2.1).

This suggests finding, for each instance x_i , the set of y 's such that $p_y(x) \geq f(\tau) \stackrel{\text{def}}{=} \frac{1}{m+2} + \tau$, where τ is uniformly distributed in $[0.5\delta, \delta]$. Specifically:

1. Select τ uniformly in $[0.5\delta, \delta]$. Indeed, selecting τ in $\{(q+i-0.5) \cdot \delta/2q : i \in [q]\}$, where $q = \text{poly}(tm)$, will do.
2. For each $i \in [t]$, taking a sample of size $\text{poly}(q/\epsilon)$, let C'_i denote the set of y 's that appeared with frequency at least $f(\tau)$ in the sample.
(Indeed, $f(\tau) \in [\frac{1}{m+2} + (\delta/2), \frac{1}{m+2} + \delta]$ is used as a random threshold for distinguishing between some valid solutions in C_x and all invalid solutions (i.e., $\{0, 1\}^* \setminus R(x) \subseteq \{0, 1\}^* \setminus C_x$).
3. For each $i \in [t]$, output the lex-first element of C'_i .

With very high probability, each C'_i is non-empty and contains only elements of C_{x_i} .¹¹ Furthermore, each $y \in C'_i$ occurs in the sample with frequency that is very close to $p_y(x_i)$. We call τ **bad** for i if for some $y \in C_{x_i}$ (or rather y such that $p_y(x_i) > 1/(m+2) = f(0)$) it holds that $|p_y(x_i) - f(\tau)| < \delta/4q$,

¹¹Actually, with very high probability, C'_i does not contain elements of $\{y : p_y(x) \leq \frac{1}{m+2}\}$ (cf., Observation (ii) in the proof of Proposition 2.2).

and observe that the probability that τ is bad for some i is $o(1/tm)$. We complete the proof by showing that the algorithm's output, when τ is not bad for any i , is one of $t \cdot (m-1) + 1$ possibilities.

To see this, fix a τ that is not bad any i , and consider the set $\{(i, y) : p_y(x_i) > f(0)\}$. Arrange all these $\leq t \cdot m$ pairs according to the value of $p_y(x_i)$, and note that $f(\tau)$ is sufficiently far from each of these values. Hence, with probability $1 - o(1/tm)$, each set C'_i equals the set of y 's such that $p_y(x_i) > f(\tau)$. It follows that, in this case, the output of the algorithm is uniquely determined by the set $\{(i, y) : p_y(x_i) > f(\tau)\}$, or, equivalently, by the location of τ in the sorted sequence of $p_y(x_i)$'s (where $i \in [t]$ and $y \in C_{x_i}$). Evidently, there are at most $tm + 1$ possible locations, but a closer examination reveal that only $tm + 1 - t = (m-1)t + 1$ are actually possible, since for each i there exists y such that $p_y(x_i) \geq f(\delta)$ (and $f(\delta) > f(\tau)$ always holds). The claim follows. ■

Adaptive invocations. We note that a result of the flavor of Theorem 3.3 for adaptive invocations of 2-pseudodeterministic algorithms (i.e., invoking the algorithm on instances determined by prior invocations) would imply efficient poly-pseudodeterministic algorithms for all “search problems in BPP” (as defined in [4, Sec. 3.1]). This is the case because “BPP search problems” are deterministically reducible (in polynomial-time) to promise problems in BPP (see [4, Thm. 3.5]), which in turn have polynomial-time 2-pseudodeterministic algorithms. (Needless to say, these reductions are adaptive.)¹²

4 A Complexity Theoretic Perspective

In this section we consider polynomial-time procedures. Confining ourselves to search problems having “efficiently recognizable solutions” (as defined next), we characterize the class of search problems that are solvable by m -pseudodeterministic (polynomial-time) algorithms in terms of (deterministic polynomial-time) reducibility to decisional problems in “promise-BPP”.

Definition 4.1 (search problems with efficiently recognizable solutions):¹³ *A search problem $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$ is said to have efficiently recognizable solutions if there exists a probabilistic polynomial time algorithm for deciding membership in R (i.e., R , as a set, is in \mathcal{BPP}).*

Our starting point is the characterization provided by Gal and Goldwasser [2].

Theorem 4.2 (characterization of the class of search problems that are solvable by polynomial-time pseudodeterministic algorithms [2]):¹⁴ *Let R be a search problem having efficiently recognizable solutions. The problem R can be solved by a polynomial-time pseudodeterministic algorithm if and only if it is reducible in deterministic polynomial-time to a decision problem in \mathcal{BPP} .*

Indeed, the reductions referred to in Theorem 4.2 are deterministic polynomial-time oracle machines that make multiple queries to the (binary) oracle, which in turn is (the characteristic function of) a set in \mathcal{BPP} . Interestingly, using reductions to the promise problem version of \mathcal{BPP} , denoted

¹²Added in Revision: The foregoing claim extends to PPT-search problems as recently defined in [5, Def. 2.4], because they are also deterministically reducible (in polynomial-time) to promise problems in BPP (see [5, Thm. 2.5]).

¹³Added in Revision: A more general notion of search problems with efficiently recognizable solutions has been recently suggested in [5]. The results presented in this section can be extended to it.

¹⁴Actually, the equivalence holds (and is stated in [2]) also for search problems not having efficiently recognizable solutions.

prBPP , allows such reductions to solve any “BPP search problem” (i.e., a search problem that can be solved in probabilistic polynomial-time and has efficiently recognizable solutions).

Recall that a **promise problem** is a pair of disjoint sets, $(S_{\text{YES}}, S_{\text{NO}})$, and solving it means distinguishing inputs in S_{YES} from inputs in S_{NO} ; the set $S_{\text{YES}} \cup S_{\text{NO}}$ is called the **promise**. Consequently, when a reduction to a promise problem makes a query that lies outside the promise (or “violates the promise”), it obtains an arbitrary answer (see [3, Sec. 2.4.1.1]).

Theorem 4.3 (characterization of the class of search problems that are solvable by probabilistic polynomial-time algorithms [4]): *Let R be a search problem having efficiently recognizable solutions. The problem R can be solved by a probabilistic polynomial-time algorithm if and only if it is reducible in deterministic polynomial-time to a promise problem in prBPP .*

Before proceeding, we highlight the fact that the syntactic difference between pure decision problems and promise problems is translated to a fundamental difference between two natural notions of probabilistic polynomial-time algorithms: pure decision problems (i.e., having a trivial promise that contains all strings)¹⁵ correspond to pseudodeterministic solvers (see Theorem 4.2), whereas promise problems correspond to general solvers (see Theorem 4.3).¹⁶

Relating prBPP and 2-pseudodeterministic algorithms. For every promise problem $\Pi = (S_{\text{YES}}, S_{\text{NO}})$, we consider the corresponding search problem $R_{\Pi} \subseteq \{0, 1\}^* \times \{0, 1\}$ such that

$$R_{\Pi}(x) = \begin{cases} \{1\} & \text{if } x \in S_{\text{YES}} \\ \{0\} & \text{if } x \in S_{\text{NO}} \\ \{0, 1\} & \text{otherwise (i.e., } x \notin S_{\text{YES}} \cup S_{\text{NO}}\text{)}. \end{cases}$$

Note that $\Pi \in \text{prBPP}$ if and only if R_{Π} can be solved in probabilistic polynomial-time. Furthermore, by employing error-reduction (which is possible under the current circumstances), such a solver can be made a 2-pseudodeterministic algorithm. Hence, Theorem 4.3 can be restated as asserting that *every search problem having efficiently recognizable solutions is reducible in deterministic polynomial-time to a search problem that can be solved by a polynomial-time 2-pseudodeterministic algorithm*. Consequently, we get the following corollary.

Theorem 4.4 (2-pseudodeterminism versus pseudodeterminism in the context of probabilistic polynomial-time): *Suppose that every search problem that can be solve by a polynomial-time 2-pseudodeterministic algorithm can also be solved by a polynomial-time pseudodeterministic algorithm. Then, every search problem that has efficiently recognizable solutions and can be solved by a probabilistic polynomial-time algorithm has a polynomial-time pseudodeterministic solver.*

Interpolating Theorems 4.2 and 4.3.

It turns out that Theorems 4.2 and 4.3 are special cases of a general result, which refers to the number of queries outside the promise that appear in the “tree of all possible executions” of the reduction (on a fixed input). The latter notion is defined next.

¹⁵Hence, the (pure) decision problem associated with the set S corresponds to the promise problem $(S, \{0, 1\}^* \setminus S)$.

¹⁶Added in Revision: This correspondence was strengthened in [1, 7].

Definition 4.5 (the directed tree of all possible executions of a reduction): *Let M be deterministic reduction to a promise problem $\Pi = (S_{\text{YES}}, S_{\text{NO}})$; that is, M is a deterministic oracle machine that makes oracle calls to Π . For any input x , the tree of all possible executions of $M^\Pi(x)$ describes all possible executions of M on input x and oracle Π such that internal nodes in this directed tree are associated with queries that may be made in such an execution, outgoing edges in the tree correspond to possible answers to these queries, and leaves correspond to the outputs in such executions. Specifically, a query q that satisfies the promise of Π has a single child in the tree, since the answer to q is 1 if $q \in S_{\text{YES}}$ and 0 if $q \in S_{\text{NO}}$, whereas a query that violates the promise has two children corresponding to the two possible answers to the query $q \notin S_{\text{YES}} \cup S_{\text{NO}}$.*

Note that (standard) decision problems have trivial promises (i.e., $S_{\text{YES}} \cup S_{\text{NO}} = \{0, 1\}^*$), and so the tree of all possible executions of a (deterministic) reduction to a decision problem consists of a single path (i.e., all internal nodes have degree 1 and the directed tree has a single leaf). On the other hand, the tree of all possible executions of a deterministic polynomial-time reduction to a promise problem may contain an exponential number of internal nodes. It turns out that the number of leaves in the tree corresponding to such a reduction equals the “amount of multi-pseudodeterminism” of solvers for the search problem (that is being reduced to Π).

Theorem 4.6 (characterization of the class of search problems that are solvable by polynomial-time m -pseudodeterministic algorithms):¹⁷ *Let R be a search problem having efficiently recognizable solutions, and $m : \mathbb{N} \rightarrow \mathbb{N}$ be upper-bounded by a polynomial. The problem R can be solved by a polynomial-time m -pseudodeterministic algorithm if and only if it is reducible in deterministic polynomial-time to a promise problem Π in prBPP such that on input x the tree of all possible executions of $M^\Pi(x)$ contains at most $m(|x|)$ leaves.*

Note that Theorem 4.2 is a special case of Theorem 4.6 in which $m \equiv 1$, whereas Theorem 4.3 can be restated by replacing the two occurrences of m (in the assertion of Theorem 4.6) by two different functions of the form $\exp(\text{poly})$.¹⁸

Proof: As a warm-up let us briefly recall the proof of Theorem 4.3. When going from a reduction to an algorithm we just emulate oracle access to a promise problem in prBPP . When $m = 1$ this yields a polynomial-time pseudorandom algorithm, since all queries satisfy the promise. Here we have to be more careful about the emulation so that $m - 1$ queries that violate the promise yield m rather than $\exp(m)$ possible outcomes. Indeed, we shall use the key idea that underlies proofs of Claim 3.2 and Theorem 3.3.

In the opposite direction (i.e., from an algorithm to a reduction) the proof of Theorem 4.3 is based on extending a prefix of a solution bit by bit. This iterative process relies on the ability to estimate the probability that the algorithm outputs a solution that fits a given prefix. Here we observe that an internal node of degree 2 is created in the tree of possible reduction only if there are two different solutions that extend the current prefix and are output by the algorithm with probability greater than $1/(m+2)$.

¹⁷Actually, the equivalence holds also for search problems not having efficiently recognizable solutions; indeed, the following proof does not refer to this hypothesis.

¹⁸The resulting assertion reads: *For any search problem R having efficiently recognizable solutions, R can be solved by a polynomial-time $\exp(\text{poly})$ -pseudodeterministic algorithm if and only if it is reducible in deterministic polynomial-time to a promise problem Π in prBPP such that on input x the tree of all possible executions of $M^\Pi(x)$ contains at most $\exp(\text{poly}(|x|))$ leaves.* This assertion is merely a cumbersome form of Theorem 4.3, since any probabilistic polynomial-time algorithm is $\exp(\text{poly})$ -pseudodeterministic and the tree of execution of any deterministic polynomial-time reduction has at most $\exp(\text{poly}(|x|))$ leaves.

From a reduction to an algorithm. Suppose that R can be solved by a deterministic (polynomial-time) oracle machine M when given oracle access to a promise problem $\Pi \in \text{pr}\mathcal{BPP}$ such that the tree of all possible executions of $M^\Pi(x)$ contains at most $m = m(|x|)$ leaves. Then, this tree contains at most $m - 1$ internal nodes that correspond to queries that lie outside the promise of Π (hereafter referred to as violating the promise). The idea is to emulate a possible execution of $M^\Pi(x)$ by replacing the oracle calls with an actual probabilistic polynomial-time computation. When the query lies inside the promise (i.e., satisfies the promise), doing so is straightforward (since by the hypothesis $\Pi \in \text{pr}\mathcal{BPP}$), and the answer is uniquely determined and is correctly obtained (with overwhelmingly high probability). However, when the query violates promise, we can estimate the acceptance probability of the machine deciding Π , and decide according to a single randomly selected threshold (as in the proofs of Claim 3.2 and Theorem 3.3). We stress that although we cannot tell (during the emulation) whether or not a query satisfies the promise, in the analysis we shall rely on the fact that the set of queries (in the tree) that violate the promise is fixed. Furthermore, any answer we provide to queries that violate the promise is OK, but we want to upper-bound the number of possible answer sequences, because this number upper-bounds the “amount of multi-pseudodeterminism”. (Note that we cannot just answer these violating queries in a fixed manner, since we do not know which queries violate the promise.)

Indeed, we (cannot and) do not need to know whether or not the query satisfies the promise; we rather estimate the acceptance probability of the machine deciding Π , and act accordingly (while relying on the fact that if the query satisfies the promise then its acceptance probability is bounded away from the selected threshold). Before detailing the resulting algorithm, we observe that all queries in the tree (including the $m - 1$ queries that violate the promise) are fully determined by M, x and Π . This implies that our decision is uniquely determined whenever the random threshold is sufficiently far from the accepting probabilities of all queries (including the $m - 1$ queries that violate the promise). As in the proof of Claim 3.2, it follows that the resulting (polynomial-time) algorithm is a $((m - 1) + 1)$ -pseudodeterministic algorithm that solves R . Specifically, using a randomized decision procedure P for Π , on input x , our randomized algorithm proceeds as follows.

1. Select a threshold $\tau \in \left\{0.4 + \frac{0.2 \cdot i}{(m+2)^2} : i \in [(m+2)^2]\right\}$ uniformly at random.
2. Emulate the execution of M on input x , answering queries by estimating the acceptance probability of P and returning 1 if and only if the estimate is larger than τ . Specifically, when M makes a query q , we estimate $p_q \stackrel{\text{def}}{=} \Pr[P(q)=1]$ such that, with probability at least $1 - o(1/m^2d)$, the estimate \tilde{p}_q falls inside $[p_q \pm 0.01 \cdot (m+2)^{-2}]$, where d is the depth of the tree of all possible executions of $M^\Pi(x)$. We return 1 if $\tilde{p}_q > \tau$ and 0 otherwise.

Note that the foregoing approximation can be obtained by $O(m^4 \log(md)) \leq \text{poly}(|x|)$ invocations of P , where the inequality uses the upper bounds on m and d (where d is upper-bounded by the running time of M).

3. When M halts, we output its verdict.

Observe that τ is very far from the p_q -value of any query that satisfies the promise of Π (i.e., for every such query q it holds that either $p_q \geq 2/3$ or $p_q \leq 1/3$, whereas $\tau \in [0.4, 0.6]$). As for queries that violate the promise of Π , with probability at most $m \cdot (m+2)^{-2} < (m+2)^{-1} - (m+2)^{-2}$, the p_q -value of some of these queries is $0.01 \cdot (m+2)^{-2}$ -close to τ . Furthermore, when the p_q -values of all queries are $0.01 \cdot (m+2)^{-2}$ -far from τ , with probability at least $1 - o(1/m^2)$, the relation of each

of the \tilde{p}_q 's to τ is correctly determined (by the relation of the p_q 's to τ). Hence, with probability at least $1 - (m + 2)^{-1}$, the foregoing algorithm answers all queries according to the relation of the p_q -values to τ (i.e., the answer to query q is 1 if and only if $p_q > \tau$). As in the proof of Claim 3.2 (see also the proof of Theorem 3.3), it follows that the foregoing algorithm is an m -pseudodeterministic solver of R .

From an algorithm to a reduction. Turning to the opposite direction, let A be an m -pseudodeterministic algorithm that solves R . The oracle machine that we construct, iteratively extends a prefix of a single solution that is output by $A(x)$ with probability at least $\frac{1+(1/m(|x|))}{m(|x|)+2}$, provided $x \in S_R$. Indeed, in each iteration the current prefix is extended by a single bit such that the probability of outputting a single solution that extends this prefix is essentially preserved (or even increased). In order to determine this bit, we need to estimate the probability that $A(x)$ outputs a single solution that fits the extended prefix, which can be done by issuing an adequate query to $\text{pr}\mathcal{BPP}$.

Note that we cannot hope distinguish in $\text{pr}\mathcal{BPP}$ between the case that the extended prefix fits a single solution that is output with probability at least p and the case that each fitting solution is output with probability smaller than p , but we can distinguish (in $\text{pr}\mathcal{BPP}$) between the former case and the case that the latter probability is smaller than $p - 1/\text{poly}(|x|)$. Hence, in the i^{th} iteration we seek a i -bit long prefix that fits some single solution that is output by $A(x)$ with probability at least $\frac{1+(1/m(|x|))-(i/\text{poly}(|x|))}{m(|x|)+2}$. We enter this iteration with a $(i-1)$ -bit long prefix y of a single solution that is output by $A(x)$ with probability at least $\frac{1+(1/m(|x|))-(i-1)/\text{poly}(|x|)}{m(|x|)+2}$, and use the oracle in order to distinguish the case that $y0$ is output by $A(x)$ with probability at least $\frac{1+(1/m(|x|))-(i-1)/\text{poly}(|x|)}{m(|x|)+2}$ and the case that $y0$ is output by $A(x)$ with probability smaller than $\frac{1+(1/m(|x|))-(i/\text{poly}(|x|))}{m(|x|)+2}$. (Specifically, if $(x, y0)$ is not a YES-instance, then $y1$ must be the extension of y that fits a solution output by $A(x)$ with probability at least $\frac{1+(1/m(|x|))-(i-1)/\text{poly}(|x|)}{m(|x|)+2}$.)

The key observation is that this query (i.e. $y0$) violates the corresponding promise if and only if both possible setting of the current bit yield a prefix that can be extended to a single solution that is output by $A(x)$ with about the same (or higher) probability. It follows that the number of leaves in the execution tree of the oracle machine that we construct cannot exceed the number of solutions that $A(x)$ outputs with probability greater than $1/(m(|x|) + 2)$, which is upper-bounded by $m(|x|)$. The resulting oracle machine, denoted M , is detailed next, while assuming (without loss of generality)¹⁹ that all solutions in $R(x)$ have length $\ell(|x|)$, for some easy to compute function $\ell : \mathbb{N} \rightarrow \mathbb{N}$.

1. On input x , estimate $\Pr[A(x) = \perp]$, and halt outputting \perp if the estimate exceeds $1/2$. (This estimate is obtained by making the query x to the \mathcal{BPP} -set $\{x' : \Pr[A(x') = \perp] \geq 2/3\} = \{0, 1\}^* \setminus S_R$, while noting that $\Pr[A(x) = \perp] \leq 1/3$ for any $x \in S_R$.)

Otherwise (i.e., $x \in S_R$), the machine initializes $y \leftarrow \lambda$ as a prefix of some solution that is output with sufficiently high probability. Indeed, at this point, there exists $z \in \{0, 1\}^{\ell(|x|)-|y|}$ such that $\Pr[A(x) = yz] \geq \frac{1+(1/m(|x|))}{m(|x|)+2}$ (see Part 1 of Proposition 2.2).

Let $\delta(n) = 1/m(n)$ and $\epsilon(n) = \delta(n)/2\ell(n)$.

¹⁹We may set ℓ to be the running time of A , and consider padding all solutions to length $\ell(|x|) + 1$; for example, consider $R'(x) \stackrel{\text{def}}{=} \{y10^{\ell(|x|)-|y|} : y \in R(x)\}$.

2. For $i = 1, \dots, \ell(|x|)$, test if extending the prefix y by 0 yields a prefix of a single solution that is output with sufficiently high probability. Specifically, we estimate the probability that $A(x)$ outputs a single solution with prefix $y0$, set $y \leftarrow y0$ if the estimate exceeds $\frac{1+\delta(|x|)-(i-1)\cdot\epsilon(|x|)}{m(|x|)+2}$, and set $y \leftarrow y1$ otherwise.

The foregoing estimate is obtained by making the query $(x, y0)$ to the promise problem $\Pi = (S_{\text{YES}}, S_{\text{NO}})$ such that

$$\begin{aligned} S_{\text{YES}} &= \left\{ (x', y') : \exists z \in \{0, 1\}^{\ell(|x'|)-|y'|} \text{ s.t. } \Pr[A(x')=y'z] \geq \frac{1 + \delta(|x'|) - (|y'| - 1) \cdot \epsilon(|x'|)}{m(|x'|) + 2} \right\} \\ S_{\text{NO}} &= \left\{ (x', y') : \forall z \in \{0, 1\}^{\ell(|x'|)-|y'|} \quad \Pr[A(x')=y'z] < \frac{1 + \delta(|x'|) - |y'| \cdot \epsilon(|x'|)}{m(|x'|) + 2} \right\} \end{aligned}$$

That is, we extend $y \in \{0, 1\}^{i-1}$ to $y0$ if and only if the query $(x, y0)$ is answered positively. Hence, if $y \in \{0, 1\}^{i-1}$ is extended to $y0$, then $(x, y0) \notin S_{\text{NO}}$ must hold, which implies that $A(x)$ outputs a single solution with prefix $y0$ with probability at least $\frac{1+\delta(|x|)-i\cdot\epsilon(|x|)}{m(|x|)+2}$. On the other hand, if $y \in \{0, 1\}^{i-1}$ is extended to $y1$, then $(x, y0) \notin S_{\text{YES}}$ must hold, which implies that $A(x)$ outputs a single solution with prefix $y1$ with probability at least $\frac{1+\delta(|x|)-(i-1)\cdot\epsilon(|x|)}{m(|x|)+2}$, since otherwise $A(x)$ outputs each single solution with prefix y with probability less than $\frac{1+\delta(|x|)-(i-1)\cdot\epsilon(|x|)}{m(|x|)+2}$, which contradicts the guarantee of the prior iteration. In both cases, the extended i -bit long y satisfies the following: there exists $z \in \{0, 1\}^{\ell(|x|)-i}$ such that $\Pr[A(x)=yz] \geq \frac{1+\delta(|x|)-i\cdot\epsilon(|x|)}{m(|x|)+2}$.

Note that the promise problem Π is in prBPP by virtue of an algorithm that, on input (x', y') , invokes $A(x')$ for $O(m(|x'|)/\epsilon(|x'|))^2 = \text{poly}(|x'|)$ times and estimate the maximal probability that the output equals a single string that extends y' .

3. Output $y \in \{0, 1\}^{\ell(|x|)}$, while noting that

$$\Pr[A(x)=y] \geq \frac{1 + \delta(|x|) - |y| \cdot \epsilon(|x|)}{m(|x|) + 2} = \frac{1 + (1/2m(|x|))}{m(|x|) + 2}$$

Hence, $y \in R(x)$ (see Part 1 of Proposition 2.2).

Note that whenever the query $(x, y0)$ lies outside $S_{\text{YES}} \cup S_{\text{NO}}$ it is the case that for both $\sigma \in \{0, 1\}$ it holds that $A(x)$ outputs a single solution with prefix $y\sigma$ with probability at least $\frac{1+\delta(|x|)-|y\sigma|\cdot\epsilon(|x|)}{m(|x|)+2}$, which equals $\frac{1+(1/m(|x|))-(|y\sigma|/2m(|x|)\ell(|x|))}{m(|x|)+2} \geq \frac{1+(1/m(|x|))-(1/2m(|x|))}{m(|x|)+2}$. Hence, nodes in the tree of all possible executions of $M^\Pi(x)$ that have two children correspond to a pair of prefixes $(y0, y1)$ that are each a prefix of a single solution output by $A(x)$ with probability exceeding $1/(m(|x|) + 2)$, which implies that these two different solutions must be in the set C_x (as in Definition 2.1).²⁰ Thus, the number of different nodes in level i of the tree is upper-bounded by the number of different i -bit long prefixes of (different) solutions in C_x . It follows that the number of leaves in this trees is at most $|C_x| \leq m(|x|)$. ■

²⁰Recall that $\Pr[A(x) \in C_x] \geq \frac{m(|x|)+1}{m(|x|)+2}$.

5 Acknowledgements and the Story Behind this Paper

I wish to thank Salil Vadhan and Ofer Grossman for their contributions to this work. In particular, the definition of m -pseudodeterministic list solver (i.e., Definition 2.5) was suggested to me by Salil Vadhan, and Part 2 of Theorem 2.10 was suggested to me by Ofer Grossman. I am grateful to both of them for the permissions to include their contributions here.

In the Fall of 2018, a group of former students of Shafi started planning to hold a small celebration of her birthday (November 14th). One day after November 14th 2018, I heard Shafi present an overview of pseudodeterminism, a direction of research she initiated in her paper with Gat [2]. It has occurred to me that nothing could be more appropriate as a gift to Shafi than trying to contribute to this research direction. The results of my attempt were first presented in the aforementioned celebration, which took place at the Simons Institute on January 13th 2019. The suggestions of Salil Vadhan and Ofer Grossman followed that presentation.

The dedication is meant to evoke the dedication of Beethoven's third ("Eroica") symphony, but (unlike in the original) avoiding Shafi's name is not meant to express dismay.

References

- [1] P. Dixon, A. Pavan, and N.V. Vinodchandran. Promise Problems Meet Pseudodeterminism. *ECCC*, TR21-043, 2021.
- [2] E. Gat and S. Goldwasser. Probabilistic Search Algorithms with Unique Answers and Their Cryptographic Applications. In *ECCC*, TR11-136, 2011.
- [3] O. Goldreich. *Computational Complexity: A Conceptual Perspective*. Cambridge University Press, 2008.
- [4] O. Goldreich. In a World of $P=BPP$. In *Studies in Complexity and Cryptography*, LNCS Vol. 6650, Springer, pages 191–232, 2011.
- [5] O. Goldreich. On Defining PPT-search Problems. This volume.
- [6] O. Goldreich, S. Goldwasser, and D. Ron. On the possibilities and limitations of pseudodeterministic algorithms. In *4th ITCS*, pages 127–138, 2013.
- [7] O. Goldreich and R. Tell. Complexity theoretic implications of pseudodeterministic algorithms for PPT-search problems. In *ECCC*, TR25-xxx, 2025.
- [8] S. Goldwasser and O. Grossman. Bipartite Perfect Matching in Pseudo-Deterministic NC. In *44th ICALP*, pages 87:1–87:13, 2017.
- [9] O. Grossman and Y.P. Liu. Reproducibility and Pseudo-Determinism in Log-Space. In *30th SODA*, pages 606–620, 2019.
- [10] M. Saks and S. Zhou. $BP_{HSpace}(S) \subseteq DSPACE(S^{3/2})$. *JCSS*, Vol. 58 (2), pages 376–403, 1999.