

Direct Product Theorems for Randomized Query Complexity

Shalev Ben-David
Institute for Quantum Computing
University of Waterloo
shalev.b@uwaterloo.ca

Eric Blais
University of Waterloo
eric.blais@uwaterloo.ca

Abstract

We establish two new direct product theorems for the randomized query complexity of Boolean functions.

The first shows that computing n copies of a function f , even with a small success probability of γ^n , requires $\Theta(n)$ times the *maximum distributional* query complexity of f with success parameter γ . This result holds for all success parameters γ , even when γ is very close to $1/2$ or to 1. As a result, it unifies and generalizes Drucker’s direct product theorem (2012) for γ bounded away from $\frac{1}{2}$ and 1 as well as the strong direct sum theorem of Blais and Brody (2019) for $\gamma \approx 1 - 1/n$.

The second establishes a general *list decoding* direct product theorem that captures many different variants of “partial computation” tasks related to the function f^n consisting of n copies of f . Notably, our list decoding direct product theorem yields a new threshold direct product theorem and other new variants such as the labelled-threshold direct product theorem.

Both of these direct product theorems are obtained by taking a new approach. Instead of directly analyzing the query complexity of algorithms, we introduce a new measure of complexity of functions that we call *discounted score*. We show that this measure satisfies a number of useful structural properties, including tensorization, that make it particularly suitable for the study of direct product questions.

Contents

1	Introduction	3
1.1	An Optimal Direct Product Theorem	4
1.2	List-Decoding Direct Product Theorem	4
1.3	Discounted Score	6
1.4	Success-conditioned score	8
1.5	Overview of the Other Proofs	9
1.5.1	Direct Product Theorem Upper Bound	9
1.5.2	List-Decoding Direct Product Theorem	10
2	Definitions and Foundations	11
2.1	Score Measures	12
2.1.1	Examples of Score Measures	13
2.2	Complexity Measures	14
2.3	Maximum Distributional Complexity	14
2.3.1	Score Amplification Properties	15
2.4	Score-Weighted Complexity	15
2.4.1	Success-Conditioned and Score-Weighted Cost	15
2.4.2	Maximum Score-weighted Distributional Query Complexity	16
3	Proof of the Direct Product Theorem	16
3.1	Proof of the Lower Bounds	17
3.2	Proof of the Upper Bound	18
4	Proof of the Tensorization Lemma	20
5	Proof of the Equivalence Lemma	22
5.1	Proof Overview	22
5.2	The Proof	23
6	List-Decoding Direct Product Theorems	25
6.1	Preliminary Lemmas	25
6.2	Proof of the List-Decoding Direct Product Theorem	28
6.3	Proof of the Threshold Direct Product Theorem	30
6.4	Proof of the Labelled Threshold DPT	30
7	Score Amplification	31
7.1	Success Amplification by Boosting	31
7.2	Hellinger Score Amplification	33
8	Relations Between the Different Complexity Measures	35
8.1	Worst-Case and Average-Case Complexity	35
8.2	Maximum Distributional Complexity	37
8.3	Score-Weighted Complexity	40
8.3.1	Relation to Maximum Distributional Complexity	40
8.3.2	Relation to Worst-Case Query Complexity	40
8.4	Score and Discounted Score	41
	Acknowledgements	41
	References	42

1 Introduction

The n -fold direct product of a function $f: \{0, 1\}^m \rightarrow \{0, 1\}$ is the function $f^n: \{0, 1\}^{m \times n} \rightarrow \{0, 1\}^n$ defined by

$$f^n(x^{(1)}, x^{(2)}, \dots, x^{(n)}) = (f(x^{(1)}), f(x^{(2)}), \dots, f(x^{(n)}))$$

for all $x^{(1)}, \dots, x^{(n)} \in \{0, 1\}^m$. How does the complexity of computing f^n relate to the complexity of f ? This question has received a lot of attention over the years in multiple different computational models, including in the setting of randomized query complexity [IRW94; NRS98; Sha04; KŠW07; JKS10; Dru12; BB19]; let us briefly make this question more precise and present the state of the art on it.

Let $R_\gamma(f)$ denote the *worst-case randomized query complexity* of f for the success probability parameter $\gamma \in (0, 1]$. This complexity measure represents the minimum number of bits of the input that a randomized algorithm must query (in the worst case over both the input and the internal randomness of the algorithm) in order to compute the value of the function on this input correctly with probability at least γ . The randomized query complexity of f^n is bounded above by

$$R_{\gamma^n}(f^n) \leq n R_\gamma(f) \tag{1}$$

since we can always compute f^n with success probability at least γ^n by running n independent instances of an algorithm that computes f with success probability at least γ . How tight is this naïve upper bound?

In a groundbreaking result, Drucker [Dru12] showed that every function f satisfies

$$R_{\gamma^n}(f^n) \geq \Omega\left(\min\left\{\left(\gamma - \frac{1}{2}\right)^3, 1 - \gamma\right\} \cdot n R_\gamma(f)\right). \tag{2}$$

This result shows that the bound (1) is asymptotically tight in the *bounded-error regime* where γ is bounded away from $\frac{1}{2}$ and 1. But it leaves open the possibility that the same bound can be improved in the *small-error regime* where $\gamma = 1 - o(1)$ and the *small-advantage regime* where $\gamma = \frac{1}{2} + o(1)$. And as it turns out, the bound (1) can indeed be strengthened in those regimes.

Let $\overline{R}_\gamma(f)$ denote the *average-case query complexity* of f , where the cost of a randomized algorithm on some input x is measured in expectation over the internal randomness of the algorithm. With standard truncation and success amplification arguments, we can show that when n is large enough,

$$R_{\gamma^n}(f^n) \leq O\left(n \overline{R}_\gamma(f)\right). \tag{3}$$

Clearly, we always have $\overline{R}_\gamma(f) \leq R_\gamma(f)$ so this bound is asymptotically always at least as strong as (1). In the bounded-error regime, $\overline{R}_\gamma(f) = \Theta(R_\gamma(f))$ so the two bounds are asymptotically equivalent. But in the small-error regime, there are partial and total functions f for which $\overline{R}_\gamma(f) = o(R_\gamma(f))$ so (3) is asymptotically stronger in this regime. And, as a strong direct sum theorem of Blais and Brody [BB19] proved, that bound is asymptotically optimal in the special case where $\gamma > 1 - 1/n$: in this setting, every function f satisfies

$$R_{\gamma^n}(f^n) \geq \Omega\left(n \overline{R}_\gamma(f)\right). \tag{4}$$

Since this result only applies to the small-error regime, it raises the question of whether there can be a single argument that yields a direct product theorem that is asymptotically optimal in both the bounded-error and small-error regimes. And, more importantly, it leaves open the possibility that (3) can also be strengthened in the small-advantage regime.

1.1 An Optimal Direct Product Theorem

Our first main result is an optimal direct product theorem which shows that (3) can indeed be further strengthened and gives a tight characterization of $R_{\gamma^n}(f^n)$ in terms of the complexity of f . Let us first introduce the complexity measure that gives this characterization.

For a distribution μ on the domain of f , define $\bar{R}_\gamma^\mu(f)$ to be the minimum expected query cost of a randomized algorithm that succeeds with probability at least γ when *both* the expected cost and success probability are measured over the internal randomness of the algorithm *and* the choice of the input $x \sim \mu$. The *maximum distributional randomized query complexity* measure of f is defined to be $\bar{R}_\gamma(f) = \max_\mu \bar{R}_\gamma^\mu(f)$.

The maximum distributional complexity measure $\bar{R}(\cdot)$ was studied in previous work of Vereshchagin [Ver98], where he implicitly shows that this measure is asymptotically equivalent to $\bar{R}(\cdot)$ in the small-error regime and to $R(\cdot)$ in the bounded-error regime. As we show in Section 8.2, however, the maximum distributional complexity measure can be asymptotically smaller than these other two measures in the small-bias regime. Also, and most importantly, this is the complexity measure that characterizes $R_{\gamma^n}(f^n)$:

Theorem 1 (Optimal Direct Product Theorem). *For every (possibly partial) function $f: \{0, 1\}^m \rightarrow \{0, 1\}$, success probability γ in the range $\frac{1}{2} < \gamma < 1$, and parameter $n \geq 1$,*

$$R_{\gamma^n}(f^n) = \Omega\left(n \cdot \bar{R}_\gamma(f)\right). \quad (5)$$

This lower bound is best possible when n is large enough: there is a universal constant c such that for every function f , success probability γ in the range $\frac{1}{2} < \gamma < 1$, and parameter $n \geq \frac{c}{(\min\{2\gamma-1, 1-\gamma\})^2}$,

$$R_{\gamma^n}(f^n) = O\left(n \cdot \bar{R}_\gamma(f)\right). \quad (6)$$

The bounds in Theorem 1 have three different implications for the different regimes:

Bounded-error regime: In this regime, $\bar{R}_\gamma(f) = \Theta(R_\gamma(f))$ so that Theorem 1 simplifies to $R_{\gamma^n}(f^n) = \Theta(n R_\gamma(f))$, recovering Drucker’s direct product theorem [Dru12].

Small-error regime: When $\gamma = 1 - \epsilon$ for some $\epsilon = o(1)$, $\bar{R}_{1-\epsilon}(f) = \Theta(\bar{R}_{1-\epsilon}(f))$ so Theorem 1 simplifies to $R_{(1-\epsilon)^n}(f^n) = \Theta(n \bar{R}_{1-\epsilon}(f))$, recovering (and generalizing) the strong direct sum theorem of Blais and Brody [BB19].

Small-advantage regime: When $\gamma = \frac{1+\delta}{2}$ for some $\delta = o(1)$, there are functions f for which $\bar{R}_\gamma(f) = o(\bar{R}_\gamma(f))$. Therefore, the characterization in terms of maximum distributional complexity is necessary for an optimal direct product theorem; no such characterization in terms of average-case query complexity can hold.

We will discuss the maximum distributional complexity measure in more detail in Section 2.3.

1.2 List-Decoding Direct Product Theorem

Theorem 1 characterizes the complexity of computing the function f^n exactly in terms of the complexity of computing f . We can also instead consider many different other hardness amplification questions by requiring algorithms to only “partially” compute f^n . Let us mention two such examples:

Threshold Direct Product: Given a function f^n and a parameter $k \leq n$, what is the complexity of algorithms that compute the value $f(x^{(i)})$ correctly on at least k of the n inputs? More precisely, define the k -of- n threshold relation $\text{Thr}_k^n: \{0, 1\}^n \rightarrow \mathcal{P}(\{0, 1\}^n)$ to be

$$\text{Thr}_k^n(y) = \{z \in \{0, 1\}^n : |\{i \in [n] : y_i = z_i\}| \geq k\}.$$

What is the worst-case randomized query complexity of $\text{Thr}_k^n \circ f$?

Labelled Threshold Direct Product: Given a function f^n and a parameter $k \leq n$, what is the complexity of algorithms that compute the value of f correctly on k of the n inputs *when the algorithms must label the k instances it has solved*? This task is not equivalent to the threshold direct product question because in this task, no errors are allowed. We can formalize it by introducing the k -of- n labelled threshold relation $\text{Label}_k^n: \{0, 1\}^n \rightarrow \mathcal{P}(\{0, 1, *\}^n)$ where

$$\text{Label}_k^n(y) = \{z \in \{0, 1, *\}^n : |\{i \in [n] : z_i \neq *\}| = k \wedge \forall i \in [n] (z_i \neq *) \Rightarrow z_i = y_i\}.$$

What is the worst-case randomized query complexity of $\text{Label}_k^n \circ f$?

The threshold direct product question has been studied in query complexity and other computational models [IK10; Ung09; Dru12]. Meanwhile, the labelled threshold direct product has received much less attention so far, in part because analyzing the benefit that an algorithm receives from the ability to choose which instance to solve appears to require very different techniques than the ones used in the threshold direct product question. In fact, even the special case of the labelled threshold direct product question where $k = 1$ is quite interesting; it corresponds to the Choice Problem [BBKW14; MT18] and it is closely related to the Correlated Samples problem introduced by Basilakis et al. [BDG+20].

These two variants of the threshold direct product relation and other natural “partial computation” questions are captured by a general direct product question related to the list-decoding task first considered in the context of error-correcting codes. For any parameters n and $\ell \leq n$, define the *list-decoding* relation $\text{LD}_\ell^n: \{0, 1\}^n \rightarrow \mathcal{P}(\binom{\{0, 1\}^n}{2^{n-\ell}})$ to be

$$\text{LD}_\ell^n(y) = \{S \subseteq \{0, 1\}^n : |S| = 2^{n-\ell} \wedge y \in S\}.$$

An algorithm correctly computes $\text{LD}_\ell^n \circ f$ on some input $x \in \{0, 1\}^{m \times n}$ when it outputs a list of $2^{n-\ell}$ vectors in $\{0, 1\}^n$ that includes the vector $(f(x^{(1)}), \dots, f(x^{(n)}))$.

The randomized query complexity of $\text{LD}_\ell^n \circ f$ is bounded above by

$$R_{\gamma^\ell}(\text{LD}_\ell^n \circ f) \leq \ell R_\gamma(f)$$

since we can always compute the value of the first ℓ instances of f with success probability γ each. We show that this upper bound is asymptotically tight in the bounded-error regime. Writing $R(f) = R_{\frac{2}{3}}(f)$ to denote the randomized query complexity of f in the bounded-error regime, we obtain the following lower bound.

Theorem 2 (List-Decoding Direct Product Theorem). *There exists a universal constant $c \in (\frac{1}{2}, 1)$ such that for every (possibly partial) function $f: \{0, 1\}^m \rightarrow \{0, 1\}$, $n \geq 1$, $\ell \in \{1, 2, \dots, n\}$, and $\gamma \in [c, 1)$,*

$$R_{\gamma^\ell}(\text{LD}_\ell^n \circ f) = \Omega(\ell R(f)).$$

Theorem 2 immediately gives a strong threshold direct product theorem.

Corollary 3 (Threshold Direct Product Theorem). *There exists a universal constant $c \in (\frac{1}{2}, 1)$ such that for every (possibly partial) function $f: \{0, 1\}^m \rightarrow \{0, 1\}$, $n \in \mathbb{N}$, $k \in \{\frac{n}{2} + 1, \dots, n\}$, and $\gamma \in [c, 1)$, if we let*

$$\tau_{n,k} = n - \log |\{x \in \{0, 1\}^n : \|x\| \leq n - k\}| = \log \left(\frac{2^n}{\sum_{i=0}^{n-k} \binom{n}{i}} \right),$$

then

$$R_{\gamma^{\tau_{n,k}}}(\text{Thr}_k^n \circ f) = \Omega(\tau_{n,k} R(f)).$$

In particular, [Corollary 3](#) implies that for every constant $\delta > 0$, there is a constant $c_2 = c_2(\delta)$ such that for all $\gamma > c_2$, the complexity of computing more than a $\frac{1}{2} + \delta$ fraction of the n instances correctly with even exponentially small success probability is bounded below by

$$R_{\gamma^n}(\text{Thr}_{(\frac{1}{2}+\delta)n}^n \circ f) = \Omega(n R(f)).$$

This result also follows from Drucker’s threshold direct product theorem [\[Dru12\]](#). But [Corollary 3](#) also applies to settings where we require the algorithm to compute only “slightly” more than $\frac{1}{2}$ of the instances correctly. Namely, it implies that for every $1 \leq t \leq \sqrt{n}$, there is constant $c_3 = c_3(t, n)$ such that for all $\gamma > c_3$,

$$R_{\gamma^{t^2}}(\text{Thr}_{\frac{n}{2}+t\sqrt{n}}^n \circ f) = \Omega(t^2 R(f)).$$

For example, with $t(n) = \sqrt{\log(n)}$, we obtain the conclusion that even computing $\frac{n}{2} + \sqrt{n \log n}$ of the n instances of f correctly with even just a polynomially-small success probability requires query complexity $\Omega(\log n \cdot R(f))$.

[Theorem 2](#) also immediately gives a direct product theorem for the Labelled Threshold problem.

Corollary 4 (Labelled Threshold Direct Product Theorem). *There exists a universal constant $c \in (\frac{1}{2}, 1)$ such that for every (possibly partial) function $f: \{0, 1\}^m \rightarrow \{0, 1\}$, $n \in \mathbb{N}$, $k \leq n$, and $\gamma \in [c, 1)$,*

$$R_{\gamma^k}(\text{Label}_k^n \circ f) = \Omega(k R(f)).$$

This result shows that computing any constant fraction $\alpha > 0$ of the n instances of f with exponentially small success probability $\gamma^{\alpha n}$ still has cost $\Omega(n R(f))$. In the other extreme setting, with $k = 1$, this result shows that the “one-of- n ” problem where the algorithm is free to pick which of the n instances it wants to solve still requires query complexity $\Omega(R(f))$.

1.3 Discounted Score

Our direct product theorems [Theorems 1](#) and [2](#) are obtained by taking a new approach. Namely, instead of trying to prove such theorems directly for randomized query complexity, we start by considering a different notion of complexity of functions.

The *discounted score* of the randomized algorithm R with respect to the function f and the distribution μ on the domain of f is

$$\text{ds}_{f,\alpha}^\mu(R) = \mathbb{E}_{\substack{D \sim R \\ x \sim \mu}} \left[\text{score}_f(D(x)) \cdot e^{-\alpha \cdot \text{cost}(D(x))} \right],$$

where the parameter α is some (tuneable) *discount factor*, $D(x)$ represents the leaf of D reached by input x , and for now we can think of $\text{score}_f(D(x))$ as being the indicator function of whether D outputs the correct value $f(x)$ on input x . The *maximum discounted score* of f with respect to μ is

$$\text{DS}_\alpha^\mu(f) = \max_R \text{ds}_{f,\alpha}^\mu(R).$$

Note that the discounted score is fundamentally different from standard measures of complexity in that it integrates the cost and the score at a “local” per-input level and maximizes this value over all randomized algorithms. This is in contrast to the usual measures of complexity where one of the cost or score is fixed as a constraint, and then the complexity measure seeks to optimize the other parameter among all the algorithms that satisfy the constraint.

Tangential remarks. The name *discounted score* is chosen to highlight the conceptual similarity between this measure and the notion of discounted rate in economics. In this setting, the discounted rate quantifies the idea that \$100 now does not have the same value as \$100 received some time in the future (e.g., because \$100 received now could be invested to accrue some interest during the next year). As such, it is natural that the discount on the value should be exponential in the delay in receiving this value. Similarly, if we think of the cost of an algorithm as the time we wait before an answer, the discounted score represents the notion that receiving a correct answer to the value of $f(x)$ now is better (or more valuable) than receiving this answer after waiting for a long time.

We do not use the analogy between discounted score and discounted rate in our arguments, but it is interesting to note that previous work on randomized query complexity have used a ratio between cost and score [BB23; BBGM22]. This approach is equivalent to discounting by $1/\text{cost}$ instead of $\exp(-\text{cost})$, which is analogous to the “hyperbolic discounting fallacy” in behavioral economics. (See, e.g., [Ain12] and the references therein.) In other words, our results can be interpreted informally as support for the observation that if one uses the economist-endorsed exponential discounting instead of the behaviorally-common hyperbolic discounting, one can get better direct product theorems for randomized query complexity.

The notion of discounted score also appears to be closely related to other well-studied notions in different fields as well. In particular, is also closely related to the *attenuation* of signals in physics. Here the notion of cost can be associated with the distance of the leaf to the root of a tree and the discounted score represents the strength of a signal from the leaf at the root of the tree when all edges have the same length. Once again, we do not use the analogy directly in our arguments but perhaps these analogies suggest that further study of discounted score measures may find other useful applications in computer science beyond the ones presented here.

Motivation. Our study of the notion of discounted score is not motivated by either of the analogies discussed above but rather by the fact that it arises naturally when we consider direct product questions for randomized algorithms. Consider the naïve upper bound (1) for computing f^n by running n independent instances of a randomized algorithm for f . The success probability of the resulting algorithm behaves multiplicatively—it is the n th power of the original success probability—whereas the cost of running both algorithms behaves additively—it is only n times the cost of the original algorithm.

The discrepancy between the behavior of the score and cost parameters appears to be problematic in the study of direct product questions. Intuitively, it seems like the analysis would be much easier in some cases if both parameters behaved identically (i.e., both were multiplicative or both were additive). But we can make both parameters behave in the same way if we consider the exponent of the cost instead of the cost itself. (We would also get this identical behavior if we instead considered the logarithm of the score instead of the score itself, though this approach doesn’t seem quite as natural when the score is an indicator function for correctness).

Tensorization. The intuition that motivated the study of discounted score for direct product questions is justified: the measure of discounted score it is particularly well suited to the study of

product functions. In fact, it satisfies a “perfect” direct product theorem property or, to say this slightly differently, it *tensorizes*.

Lemma 5 (Tensorization Lemma). *For any (possibly partial) function $f: \{0,1\}^m \rightarrow \{0,1\}$, any distribution μ on $\{0,1\}^m$, and any parameter $n \geq 1$,*

$$\text{DS}_\alpha^{\mu^n}(f^n) = \text{DS}_\alpha^\mu(f)^n.$$

We prove Lemma 5 establishing inequalities in both directions. The lower bound $\text{DS}_\alpha^{\mu^n}(f^n) \geq \text{DS}_\alpha^\mu(f)^n$ is easily obtained by analyzing the algorithm for f^n constructed by running independent instances of an algorithm A for f with optimal discounted score on each of the n inputs.

The upper bound $\text{DS}_\alpha^{\mu^n}(f^n) \leq \text{DS}_\alpha^\mu(f)^n$ is more intricate. It uses an embedding argument, where we use an algorithm B for f^n that has large discounted score to solve f by embedding the input to f as the i th instance to f^n for a randomly chosen $i \in [n]$ and simulating B on this embedded input and $n - 1$ other “fake” inputs. The embedding argument in the proof of Lemma 5 departs from previous embedding arguments in that the fake inputs are not generated statically (according to μ or some other cleverly chosen distribution). Instead, we simulate B on a query-by-query basis. Whenever B queries a bit of the real input, we make the corresponding query to provide its answer. And when B queries one of the fake inputs, we generate the answer to this query according to some distribution ν conditioned on the bits that we have observed so far. By choosing ν to be the μ^n distribution conditioned on B being correct on all inputs, we obtain exactly the desired lower bound.

1.4 Success-conditioned score

In order to obtain a direct product theorem from Lemma 5, we need to relate the discounted score measure to other measures of randomized query complexity. We do so by going through another variant of randomized query complexity, which we call *success-conditioned* randomized query complexity and denote by $\overline{\text{R}}_\gamma(f)$.

Roughly speaking, $\overline{\text{R}}_\gamma(f)$ is defined similarly to the maximum distributional complexity $\overline{\text{R}}(f)$, with the one important difference that the cost of the algorithm is measured *conditioned on the algorithm succeeding*. We introduce the notion formally in Section 2.4. But for now let us instead introduce some of its important properties. First, it satisfies $\overline{\text{R}}_\gamma(f) = \Theta(\overline{\text{R}}_\gamma(f))$ over all $\gamma \in (1/2, 1)$ whenever f is a Boolean-valued function. For functions with non-Boolean output alphabets (such as f^n), however, the success-conditioned query complexity is not equivalent to $\overline{\text{R}}(f)$ and has better behavior.

Intuitively, conditioning on success helps eliminate the following type of degenerate randomized algorithm: an algorithm may attempt to solve f^n with small success probability γ^n by querying everything with probability γ^n , and querying nothing (guessing randomly instead) with probability $1 - \gamma^n$. Since the input size grows linearly with n and the desired success decreases exponentially with n , the expected cost of this algorithm is less than 1 for large n (even though the worst-case cost is very high). This means that expected cost measures such as $\overline{\text{R}}_{\gamma^n}(f^n)$ and $\overline{\text{R}}_{\gamma^n}(f^n)$ are trivial for all f (but worst-case cost measures like $\text{R}_{\gamma^n}(f^n)$ are not). Our introduction of $\overline{\text{R}}$ as a success-conditioned cost allows us to give an expected-cost version of randomized query complexity for which the direct product problem is nontrivial.

The direct product theorem is obtained from Lemma 5 by relating success-conditioned complexity to both maximum distributional complexity and maximum discounted score in the following ways.

Lemma 6 (Equivalence Lemma). *For every (possibly partial) function $f: \{0,1\}^m \rightarrow \{0,1\}$, every distribution μ over the domain of f , every success parameter γ , and every discount factor $\alpha \geq 0$,*

$$\overline{\text{R}}_\gamma^\mu(f) \geq \frac{1}{\alpha} \log \frac{\text{DS}_\alpha^\mu(f)}{\gamma}.$$

Moreover, this lower bound is best possible: for every f , μ , and γ , there exists a discount factor α^* for which

$$\overline{\overline{\text{SR}}}_\gamma^\mu(f) = O\left(\frac{1}{\alpha^*} \log \frac{\text{DS}_{\alpha^*}^\mu(f)}{\gamma}\right).$$

The tensorization and equivalence lemmas immediately give an optimal direct product theorem for maximum score-weighted distributional complexity.

Theorem 7 (Direct Product Theorem for sR). *For every (possibly partial) function $f: \{0, 1\}^m \rightarrow \{0, 1\}$, every success parameter γ in the range $\frac{1}{2} \leq \gamma \leq 1$, and every $n \geq 1$,*

$$\overline{\overline{\text{SR}}}_{\gamma^n}(f^n) = \Omega\left(n \cdot \overline{\overline{\text{SR}}}_\gamma(f)\right).$$

Note that [Theorem 7](#) is the only result of its type that we know of where we have the same randomized query complexity measure on both sides of the inequality, perfect product behavior on the success probability parameter, and no loss on the right-hand side that is a function of this parameter. As we have already discussed above, it is known that no such theorem can hold for the measures $\text{R}(\cdot)$, $\overline{\text{R}}(\cdot)$, or $\overline{\overline{\text{R}}}(\cdot)$.

The lower bound of [Theorem 1](#) is obtained as a direct corollary of [Theorem 7](#) when we observe that worst-case randomized complexity always upper bounds the success-conditioned maximum distributional complexity of any function and that for Boolean-valued functions f , the complexity measures $\overline{\overline{\text{R}}}_\gamma(f)$ and $\overline{\overline{\text{SR}}}_\gamma(f)$ can differ by at most a multiplicative factor of 2.

To summarize, the measures of randomized query complexity that we have discussed satisfy

$$\text{R}_\gamma(f) \geq \overline{\text{R}}_\gamma(f) \geq \overline{\overline{\text{SR}}}_\gamma(f) \geq \overline{\overline{\text{R}}}_\gamma(f).$$

For Boolean-valued functions, the right-most two collapse. If, additionally, γ is bounded away from $1/2$, the rightmost three collapse. Finally, if γ is also bounded away from 1, then all of these measures coincide. A direct product theorem can have $\overline{\overline{\text{SR}}}_\gamma$ on both sides (that is, measuring both f^n and f with the same notion of query complexity, $\overline{\overline{\text{SR}}}$). An even better measure, which has a perfect direct product theorem, is $\text{DS}_\alpha^\mu(f)$, but it is parametrized by a tricky-to-interpret discount factor α instead of by the success γ .

1.5 Overview of the Other Proofs

1.5.1 Direct Product Theorem Upper Bound

For the upper bound in [Theorem 1](#), our goal is to give an algorithm for computing n copies of f with success probability γ^n using a total worst-case number of queries of only $O(n \overline{\overline{\text{R}}}_\gamma(f))$. The basic idea is to run an algorithm for f on each of the n copies, but several issues arise when we try to implement this idea.

First, the measure $\overline{\overline{\text{R}}}_\gamma(f)$ does not have a minimax theorem. This means there is no algorithm which achieves such low query complexity on all distributions. Instead, a single algorithm for f can only guarantee that it achieves this expected cost and expected success probability γ on a single distribution μ . To mitigate this limitation, we can use a minimax theorem for $\text{R}_{\gamma^n}(f^n)$ instead, since it is simply worst-case query complexity. This gives us a hard distribution μ for f^n (which may not be a product distribution). We then let μ_1 be the marginal distribution of μ restricted to the first copy of f , and run an $\overline{\overline{\text{R}}}_\gamma(f)$ type algorithm for f with respect to μ_1 on the first copy. Next, we let μ_2 be the marginal for the second copy, but conditioned on the bits of the first instance

that have been revealed by the algorithm that we ran on the first copy of f . Similarly, μ_3 will be the marginal of μ on the third copy of f , conditioned on the input bits revealed by the queries of the previous two algorithms. Since each of these distribution is well-defined, we can find algorithms R_1, R_2, \dots, R_n for f with respect to these distributions, such that all of them have success at least γ and expected cost at most $\bar{R}_\gamma(f)$.

Having run these algorithms, our success probability for all copies is γ^n (against μ), as desired. Our query bound is $n\bar{R}_\gamma(f)$, also as desired. However, there is another issue: this upper bound on the number of queries only holds in expectation against μ , rather than in the worst case. To address this issue, we can impose a cutoff: terminate the algorithm if the number of queries surpasses its expectation by a factor of 10. Such a cutoff will reduce the success probability, so we want the probability that it happens to be very low, since we already only have success probability that is exponentially small in n . It turns out that when n is large enough, the runs of the algorithms on different copies of f are sufficiently independent that the total expected cost has thin tails (and a standard deviation of $O(\sqrt{n})$). From this we conclude that the success probability only decreases by the required exponentially small amount.

To finish the argument, we need one more ingredient: success amplification. We cannot amplify a generic algorithm for f^n : it places probability less than $1/2$ on the right answer, so even taking a majority vote might amplify a wrong answer instead. So we must amplify $\bar{R}_\gamma(f)$ and run the whole argument above starting with a larger value for γ . This approach leads to another complication: since \bar{R} is a distributional measure with no minimax theorem, it is not obvious that its success can be amplified. In fact, for *fixed* distributional complexity measures, success amplification does not hold. But, as we show, we do get efficient success amplification for the maximum distributional complexity measure by adapting Schapire's boosting algorithm [Sch90] to this setting.

1.5.2 List-Decoding Direct Product Theorem

For this result, we want to prove a lower bound on the query complexity of an algorithm for f^n which outputs $2^{n-\ell}$ different n -bit strings such that the probability that the correct output string for f^n is in that list is at least γ^ℓ . Fix a distribution μ over inputs to f , and consider the behavior of such an algorithm when run on inputs from μ^n . Since the distribution is fixed, we can assume that such an algorithm is deterministic, that is to say, a decision tree.

When the overall success probability of the algorithm is γ^ℓ , there must be at least a $\gamma^\ell/2$ fraction of the leaves of the decision tree for which the success probability (conditioned on reaching this leaf) is at least $\gamma^\ell/2$. For such a successful leaf L , we consider the conditional distribution of μ^n conditioned on reaching this leaf. A lemma of Drucker [Dru12] says that this conditional distribution is still a product distribution: indeed, it is just the product, over $i \in [n]$, of the conditional distribution $\mu|_{L_i}$ conditioned on the bits L_i queried to the i -th input on the path to this leaf.

Since L is a successful leaf, we know that there is a list of $2^{n-\ell}$ strings such that a sample from this product distribution $\mu|_{L_1} \times \mu|_{L_2} \times \dots \times \mu|_{L_n}$ is in the list with probability at least $\gamma^\ell/2$. Without loss of generality, the list in question can simply be the $2^{n-\ell}$ strings with highest probability in this product distribution. So a successful leaf places a significant probability mass on its top $2^{n-\ell}$ most probable strings (from the posterior distribution $\mu^n|_L$ over n -bit strings). Since the distribution in question is product, this condition can be translated into an entropy condition: the only way for such high mass to be placed on so few strings in a product distribution is for the product distribution to have fairly low entropy: entropy at most $n - \Omega(\ell)$.

Our situation now is that a decent fraction of the leaves of the decision tree (at least $\gamma^\ell/2$ of them) are such that conditioned on that leaf, the posterior distribution over the output string has entropy at most $n - \Omega(\ell)$. The next key insight is that we can use the scoring rule $2^{-H(\cdot)}$, where H is

the binary entropy function, to measure the correctness of an algorithm (or a leaf) instead of using success probability. The score we give to a leaf L of a decision tree computing f^n is the product of the scores of the n partial assignments L_i (each with respect to the distribution μ). That is, the score of L is $\prod_{i=1}^n 2^{-H(f(\mu|_{L_i}))} = 2^{-H(f^n(\mu^n|_L))}$, and the score of the decision tree is the expected score of the leaves.

It turns out that [Theorem 1](#) holds for all scoring functions which have the multiplicative property described above in which the score of a leaf in a decision tree for f^n is the product of the scores of the partial assignments L_i , together with a few other simple conditions such as non-negativity. We can therefore employ our direct product theorem with this new scoring measure $2^{-H(\cdot)}$, which plays the role of “success probability” for our purposes.

The entropy upper bound of $n - \Omega(\ell)$ on a successful leaf L translates into a score lower bound of $2^{\Omega(\ell)-n}$ at that leaf (this works because the entropy $H(\cdot)$ is additive over the independent copies of f , which translates to $2^{-H(\cdot)}$ being multiplicative over those copies). Since there are $\Omega(\gamma^\ell)$ fraction of leaves which have this score lower bound, the overall expected score of the decision tree is at least $\gamma^\ell 2^{\Omega(\ell)-n}$. The direct product theorem with this scoring rule now gives a lower bound on the number of queries used by this decision tree: it is at least $\Omega(n)$ times the cost of computing a single copy of f with a score of at least $(\gamma^\ell 2^{\Omega(\ell)-n})^{1/n}$. This latter score works out to $2^{-1+\Omega(\ell/n)}$ when $\gamma < 1$ a sufficiently large constant.

Since the score function is $2^{-H(\cdot)}$, computing f with a score of $2^{-1+\Omega(\ell/n)}$ is like attaining binary entropy of at most $1 - \Omega(\ell/n)$ for the posterior distribution over the Boolean output of f . The number of queries required to do this can be related to the number of queries to compute f with other scoring rules, such as the ones from [\[BB23\]](#). One of the scoring rules in that paper can be linearly amplified. Using this amplification property, we can take an algorithm which computes f with binary entropy $1 - \Omega(\ell/n)$, repeat it $O(n/\ell)$ times, and obtain an algorithm which computes f to bounded error. We conclude that the query cost of obtaining binary entropy $1 - \Omega(\ell/n)$ is $\Omega(R(f) \cdot \ell/n)$, at least when μ is chosen to be an appropriate hard distribution.

Putting this together with the direct product theorem, we obtain a lower bound of $\Omega(n)$ times $\Omega(R(f) \cdot \ell/n)$, which gives the desired bound of $\Omega(\ell R(f))$.

2 Definitions and Foundations

Throughout the paper, the symbol f is reserved for representing a total or partial Boolean function $f: X \rightarrow \{0, 1\}$ for some $X \subseteq \{0, 1\}^m$. The symbol μ is reserved for distributions over the domain X of f .

Since we consider only query complexity, we associate deterministic algorithms with deterministic decision trees and use both terms interchangeably. Every internal node of a decision tree is labelled with the index of a bit of the input, and the outgoing edges from these nodes are labelled with 0 and 1. These nodes correspond to the queries made by the algorithm. Every leaf of a decision tree is typically associated with a label in $\{0, 1\}$ indicating the output value of the decision tree at the leaf. (As we will see in the next section, however, in distributional settings we will ignore the value of the label of the leaves when measuring the correctness score of decision trees.)

For each leaf ℓ of a decision tree D , we use ℓ to represent both the leaf itself and the subset of inputs in X that induce a path to ℓ in D . For any input $x \in X$, we write $D(x)$ to represent the leaf in D reached by x .

A *randomized algorithm* R corresponds to a distribution over deterministic algorithms. We write R to denote both the algorithm itself and the associated distribution over deterministic decision trees. A *leaf* of R is a leaf in one of the deterministic decision trees in its support.

All the definitions carry over to the more general setting of direct product functions $f^n: X^n \rightarrow \{0, 1\}$, $X \subseteq \{0, 1\}^m$ as-is, with the single exception that the labels of leaves of decision trees for these functions are labelled with values in $\{0, 1\}^n$ or $[0, 1]^n$.

2.1 Score Measures

The score of a randomized algorithm measures how well it computes a given function f . The most common score considered in the study of randomized algorithms is success probability. Our results, however, also apply to other score measures and we will use some of these other score measures for the proof of [Theorem 2](#) so we introduce a bit of notation to allow the statement of the more general results.

Scoring measures for Boolean functions are defined in terms of a *score function* $\phi: [0, 1] \rightarrow \mathbb{R}$ that satisfies a few properties:

Boundedness: The range of ϕ is $[\frac{1}{2}, 1]$.

Normalization: The function satisfies $\phi(0) = \phi(1) = 1$ and $\phi(\frac{1}{2}) = \frac{1}{2}$.

Symmetry: For every $p \in [0, 1]$, $\phi(p) = \phi(1 - p)$.

Monotonicity: The function ϕ is monotonically increasing on $[\frac{1}{2}, 1]$.

Continuity: The function ϕ is continuous on $[0, 1]$.

The *score* of a leaf ℓ of a decision tree D with respect to a target function f and a distribution μ is defined in terms of the underlying scoring function ϕ . Recall that each leaf ℓ of D also represents a subset of the domain X of f corresponding to the set of inputs that lead to this leaf. We also write $\ell_0 = \ell \cap f^{-1}(0)$ and $\ell_1 = \ell \cap f^{-1}(1)$. We define the score of ℓ to be

$$\text{score}_f^\mu(\ell) = \phi\left(\frac{\mu(\ell_1)}{\mu(\ell)}\right).$$

Letting $\ell \sim D(\mu)$ denote the distribution on the leaves of a decision tree D induced by the distribution μ on the domain of f , the score of a randomized algorithm R with respect to f and μ is

$$\overline{\text{score}}_f^\mu(R) = \mathbb{E}_{D \sim R} \left[\mathbb{E}_{\ell \sim D(\mu)} [\text{score}_f^\mu(\ell)] \right].$$

Let us mention a few remarks about these score measures:

1. For completeness, the notation score_f^μ should also include reference to the underlying scoring function ϕ . We omit it for notational clarity and instead use different notation introduced below whenever we consider a specific score measures.
2. The score measures do not use the values of labels of decision trees when measuring their score, they can instead be thought of measures of discrepancy of the leaves with respect to f and μ .
3. The score measures can also be obtained from *scoring rules* for forecasting algorithms as introduced in [\[BB20; BB23\]](#). These scoring rules can provide additional useful structural properties to the score measures, but all our results only require the basic properties of score functions listed above so the extension to forecasting algorithms is not needed.

We also extend score measures for product functions. Consider a decision tree D that computes f^n for some function $f: \{0, 1\}^m \rightarrow \{0, 1\}$. Each leaf ℓ of D corresponds to a subset of $\{0, 1\}^{m \times n}$. We can always write such a leaf as a product $\ell = \ell^{(1)} \times \dots \times \ell^{(n)}$ where each $\ell^{(i)} \subseteq \{0, 1\}^m$. We define the score of the leaf ℓ of D with respect to a product measure μ^n on the domain of f^n to be

$$\text{score}_{f^n}^{\mu^n}(\ell) = \prod_{i=1}^n \text{score}_f^{\mu}(\ell^{(i)}).$$

As before, the score of a randomized algorithm R with respect to f^n and μ^n is the expected score of its leaves:

$$\overline{\text{score}}_{f^n}^{\mu^n}(R) = \mathbb{E}_{D \sim R} \left[\mathbb{E}_{\ell \sim D(\mu^n)} [\text{score}_{f^n}^{\mu^n}(\ell)] \right].$$

2.1.1 Examples of Score Measures

We consider the following three score measures to establish [Theorem 1](#) and [Theorem 2](#):

Success probability score. The most common way to measure the quality of a randomized algorithm is via its *success probability*. The corresponding score measure is obtained by considering the scoring function

$$\phi_{\text{success}}(p) = \max\{p, 1 - p\}.$$

The score of a randomized algorithm R with respect to μ and f when we use ϕ_{success} is

$$\overline{\text{success}}_f^{\mu}(R) = \mathbb{E}_{D \sim R} \left[\mathbb{E}_{\ell \sim D(\mu)} \max \left\{ \frac{\mu(\ell_0)}{\mu(\ell)}, \frac{\mu(\ell_1)}{\mu(\ell)} \right\} \right].$$

This is the average success probability of R over μ when each leaf of the trees in R 's support are labelled optimally. In other words, there is a labelling of the leaves of the trees in the support of R for which

$$\overline{\text{success}}_f^{\mu}(R) = \Pr_{\substack{x \sim \mu \\ D \sim R}} [D \text{ outputs } f(x) \text{ on input } x].$$

By definition, the score of a randomized algorithm R that computes f^n with respect to the product distribution μ^n is

$$\overline{\text{success}}_{f^n}^{\mu^n}(R) = \mathbb{E}_{\substack{D \sim R \\ x \sim \mu^n}} \left[\prod_{i=1}^n \text{success}_f^{\mu}(D(x)^{(i)}) \right].$$

By Drucker's product lemma (Lemma 3.2 in [\[Dru12\]](#)), this is exactly the probability that R computes $f^n(x)$ correctly when $x \sim \mu^n$:

$$\overline{\text{success}}_{f^n}^{\mu^n}(R) = \Pr_{\substack{x \sim \mu^n \\ D \sim R}} [D \text{ outputs } f^n(x) \text{ on input } x].$$

Hellinger score. We obtain a different score measure by considering the *Hellinger (squared)* scoring function

$$\phi_{\text{hel}}(p) = 1 - \sqrt{p(1 - p)}.$$

The Hellinger score of a randomized algorithm R with respect to the function f and distribution μ is

$$\overline{\text{Hel}}_f^{\mu}(R) = \mathbb{E}_{D \sim R} \left[\mathbb{E}_{\ell \sim D(\mu)} \left[1 - \sqrt{\frac{\mu(\ell_0)}{\mu(\ell)} \cdot \frac{\mu(\ell_1)}{\mu(\ell)}} \right] \right] = 1 - \mathbb{E}_{D \sim R} \sum_{\ell \in D} \sqrt{\mu(\ell_0)\mu(\ell_1)}.$$

Exponential entropy score. A third way to measure the correctness of algorithms is to consider the *exponential entropy* scoring function

$$\phi_{\text{ent}}(p) = 1 - \frac{1}{2}2^{-h(p)} = 1 - \frac{1}{2}p^p(1-p)^{1-p}$$

where $h: p \mapsto -p \log(p) - (1-p) \log(1-p)$ is the binary entropy function. The exponential entropy score of a randomized algorithm R with respect to f and μ is

$$\overline{\text{Ent}}_f^\mu(R) = 1 - \frac{1}{2} \mathbb{E}_{D \sim R} \left[\mathbb{E}_{\ell \sim D(\mu)} \mu(\ell) \cdot 2^{-h\left(\frac{\mu(\ell_0)}{\mu(\ell)}\right)} \right].$$

2.2 Complexity Measures

The *cost* of the leaf ℓ of the decision tree D , denoted $\text{cost}(\ell)$, is the depth of that leaf. This depth corresponds to the number of queries made by the algorithm before it terminates. Recall that $D(x)$ denotes the leaf of D reached by input x . The (*absolute*) *worst-case cost* of the randomized algorithm R is

$$\text{cost}(R) = \max_{\substack{x \in \text{dom}(f) \\ D \in \text{supp}(R)}} \text{cost}(D(x)).$$

The *worst-case average cost* of R is

$$\overline{\text{cost}}(R) = \max_{x \in \text{dom}(f)} \mathbb{E}_{D \sim R} [\text{cost}(D(x))].$$

The score and cost definitions introduced above give rise to a variety of different complexity measures on functions. In all of the following definitions, fix any scoring measure score_f . The *worst-case* randomized query complexity of f is

$$R_\gamma(f) = \min_{R: \text{score}_f(R) \geq \gamma} \text{cost}(R)$$

and the *average-case* complexity of f is

$$\overline{R}_\gamma(f) = \min_{R: \text{score}_f(R) \geq \gamma} \overline{\text{cost}}(R).$$

For the reader's convenience, we review some of the basic facts relating worst-case and average-case query complexity in [Section 8.1](#). Our results also rely on additional notions of randomized query complexity. We introduce these alternative measures in the next two subsections.

2.3 Maximum Distributional Complexity

The *average cost* of R over the distribution μ is

$$\overline{\text{cost}}^\mu(R) = \mathbb{E}_{x \sim \mu, D \sim R} [\text{cost}(D(x))].$$

The *distributional query complexity* of f with respect to the distribution μ and the score parameter γ is

$$\overline{R}_\gamma^\mu(f) = \min_{R: \text{score}_f^\mu(R) \geq \gamma} \overline{\text{cost}}^\mu(R),$$

the minimum expected cost of a randomized algorithm with expected score γ , where both the cost and score are taken in expectation over the internal randomness of the algorithm and the choice of the input $x \sim \mu$.

The *maximum distributional randomized query complexity* of f for the score parameter γ is

$$\overline{\overline{R}}_\gamma(f) = \max_{\mu} \overline{\overline{R}}_\gamma^\mu(f) = \max_{\mu} \min_{R: \overline{\text{score}}_\mu^\mu(R) \geq \gamma} \overline{\text{cost}}^\mu(R).$$

This measure of complexity is distribution-independent.

Note that because both the score and the cost of the randomized algorithms are taken in expectation over both the internal randomness and the choice of input over the underlying distribution, standard minimax theorems do not apply and so (unlike in the more common distributional settings considered in the literature where one of these measures is taken in a worst-case fashion) it is not clear whether $\overline{\overline{R}}_\gamma(f)$ should always equal its deterministic analogue $\overline{\overline{D}}_\gamma(f)$ or not.

2.3.1 Score Amplification Properties

The standard success amplification and boosting arguments used when considering $R(f)$ and $\overline{R}(f)$ do not apply for the maximum distributional complexity measure. Nevertheless, we show that the standard arguments can be extended to show that this complexity measure does satisfy the usual score amplification properties.

First, we show that Schapire’s boosting algorithm [Sch90] can be extended to give a success amplification theorem analogous to the ones for worst-case and average-case complexity for maximum distributional complexity.

Lemma 8. *For any $\delta > 0$, any $f: \{0, 1\}^m \rightarrow \{0, 1\}$, and any success probability parameter $\gamma \in [\frac{1+\delta}{2}, 1 - \delta]$,*

$$\overline{\overline{R}}_{\gamma + \frac{\delta}{3}}(f) \leq 4 \cdot \overline{\overline{R}}_\gamma(f).$$

We also prove a (stronger) linear score amplification lemma for the Hellinger score measure.

Lemma 9. *For any $\delta \in [0, 1]$ and any $f: \{0, 1\}^m \rightarrow \{0, 1\}$,*

$$R(f) = O\left(\frac{1}{\delta} \cdot \overline{\overline{R}}_{\text{Hel } \frac{1+\delta}{2}}(f)\right).$$

Note that these two score amplification properties do *not* hold in the standard distributional setting where we consider distributional complexity measures over a fixed distribution (c.f., [Sha04; BKST24]). The proofs of Lemmas 8 and 9 strongly rely on the fact that the maximum distributional complexity measure holds over all distributions. These proofs are presented in Section 7.

2.4 Score-Weighted Complexity

In addition to the complexity measures introduced above, we also introduce one more notion of complexity which is based on a different way to measure the cost of algorithms.

2.4.1 Success-Conditioned and Score-Weighted Cost

The average cost of a randomized algorithm R with respect to a distribution μ is measured on average over all inputs x . We can also measure the expected cost of R only on the inputs for which it outputs the correct value, giving a *success-conditioned* average cost

$$\mathbb{E}_{\substack{x \sim \mu \\ D \sim R}} [\text{cost}(D(x)) \mid D \text{ outputs } f(x)].$$

This success-conditioned cost of randomized algorithms can also be generalized to a *score-weighted cost* measure for any scoring rule by setting

$$\overline{\text{scost}}_f^\mu(R) = \frac{\mathbb{E}_{x \sim \mu, D \sim R} [\text{score}_f(D(x)) \cdot \text{cost}(D(x))]}{\overline{\text{score}}_f^\mu(R)}.$$

As we can verify directly, this measure coincides with the success-conditioned average cost when we set $\text{score}_f = \text{success}_f$ and use the optimal labels on the decision tree leaves.

The average cost and score-weighted cost of a randomized algorithm R can differ substantially. Consider for example a function $f: \{0, 1\}^m \rightarrow \{0, 1\}$ and an algorithm R for f^n that queries all mn bits of the input to determine the value of $f^n(x)$ with probability $p = \frac{1}{2^n - 1}$ and otherwise guesses the value of $f^n(x)$. For any distribution μ , the average cost of this algorithm is $\overline{\text{cost}}^\mu(R) = pmn = \frac{mn}{2^n - 1}$ but its score-weighted cost with the success score function is

$$\overline{\text{scost}}_{f^n}^\mu(R) = \frac{p \cdot mn + (1 - p)2^{-n} \cdot 0}{p + (1 - p)2^{-n}} = \frac{mn}{2} \gg \overline{\text{cost}}^\mu(R).$$

2.4.2 Maximum Score-weighted Distributional Query Complexity

The notion of score-weighted cost gives rise to one more distributional measure of complexity, the *distributional score-weighted query complexity*:

$$\overline{\text{sR}}_\gamma^\mu(f) = \min_{R: \overline{\text{score}}_f^\mu(R) \geq \gamma} \overline{\text{scost}}_f^\mu(R).$$

Once again, we obtain a distribution-free complexity measure by taking the maximum distributional score-weighted query complexity over all choices of μ :

$$\overline{\overline{\text{sR}}}_\gamma(f) = \max_\mu \overline{\text{sR}}_\gamma^\mu(f).$$

The maximum score-weighted distributional query complexity is related to both the maximum distributional and worst-case randomized query complexity measures in the following ways.

Proposition 10. *For every function $f: \{0, 1\}^m \rightarrow \{0, 1\}$ and score parameter γ ,*

$$\frac{1}{2} \cdot \overline{\text{R}}_\gamma(f) \leq \overline{\overline{\text{sR}}}_\gamma(f) \leq \frac{1}{\gamma} \cdot \overline{\text{R}}_\gamma(f)$$

Proposition 11. *For every function $f: \{0, 1\}^m \rightarrow \{0, 1\}$, parameter $n \in \mathbb{N}$, and score γ ,*

$$\overline{\overline{\text{sR}}}_{\gamma^n}(f) \leq \text{R}_{\gamma^n}(f).$$

We complete the proofs of these two propositions in [Section 8.3](#).

3 Proof of the Direct Product Theorem

We complete the proofs of the standard Direct Product Theorem, [Theorem 1](#), as well as the closely related [Theorem 7](#) in this section. We begin by establishing the lower bounds for both of these theorems.

3.1 Proof of the Lower Bounds

As discussed in the introduction, there are two main technical results at the heart of the proofs of the Direct Product Theorems. The first is the tensorization lemma for discounted score, which we restate here for the reader's convenience.

Lemma 5 (Tensorization Lemma). *For any (possibly partial) function $f: \{0,1\}^m \rightarrow \{0,1\}$, any distribution μ on $\{0,1\}^m$, and any parameter $n \geq 1$,*

$$\text{DS}_\alpha^{\mu^n}(f^n) = \text{DS}_\alpha^\mu(f)^n.$$

The second is the equivalence lemma which shows the relation between discounted score and score-weighted distributional complexity. We also restate it for convenience.

Lemma 6 (Equivalence Lemma). *For every (possibly partial) function $f: \{0,1\}^m \rightarrow \{0,1\}$, every distribution μ over the domain of f , every success parameter γ , and every discount factor $\alpha \geq 0$,*

$$\overline{\text{sR}}_\gamma^\mu(f) \geq \frac{1}{\alpha} \log \frac{\text{DS}_\alpha^\mu(f)}{\gamma}.$$

Moreover, this lower bound is best possible: for every f , μ , and γ , there exists a discount factor α^* for which

$$\overline{\text{sR}}_\gamma^\mu(f) = O\left(\frac{1}{\alpha^*} \log \frac{\text{DS}_{\alpha^*}^\mu(f)}{\gamma}\right).$$

We complete the proof of [Lemma 5](#) in [Section 4](#) and the proof of [Lemma 6](#) in [Section 5](#). But first, let us show how these lemmas let us complete the proof of the lower bound for the Direct Product Theorem for maximum score-weighted distributional complexity.

Theorem 7 (Direct Product Theorem for sR). *For every (possibly partial) function $f: \{0,1\}^m \rightarrow \{0,1\}$, every success parameter γ in the range $\frac{1}{2} \leq \gamma \leq 1$, and every $n \geq 1$,*

$$\overline{\text{sR}}_{\gamma^n}(f^n) = \Omega\left(n \cdot \overline{\text{sR}}_\gamma(f)\right).$$

Proof. Choose a distribution μ on the domain of f that maximizes $\overline{\text{sR}}_\gamma^\mu(f)$, so that $\overline{\text{sR}}_\gamma(f) = \overline{\text{sR}}_\gamma^\mu(f)$. Then if we choose α^* to be the parameter for which the second inequality in [Lemma 5](#) holds,

$$\begin{aligned} \overline{\text{sR}}_{\gamma^n}^{\mu^n}(f^n) &\geq \frac{1}{\alpha^*} \log \frac{\text{DS}_{\alpha^*}^{\mu^n}(f^n)}{\gamma^n} && \text{(First bound in [Lemma 6](#))} \\ &= \frac{1}{\alpha^*} \log \frac{\text{DS}_{\alpha^*}^\mu(f)^n}{\gamma^n} && \text{(Lemma 5)} \\ &= n \cdot \frac{1}{\alpha^*} \log \frac{\text{DS}_{\alpha^*}^\mu(f)}{\gamma} \\ &= \Omega\left(n \cdot \overline{\text{sR}}_\gamma^\mu(f)\right) && \text{(Second bound in [Lemma 6](#))} \end{aligned}$$

□

And from there we complete the proof of the lower bound for our main Direct Product Theorem. While this lower bound is stated in terms of success probability in the introduction for simplicity, we show here that it also holds for arbitrary score measures.

Theorem 12 (Lower bound of Theorem 1, generalized). *For every function $f: \{0, 1\}^m \rightarrow \{0, 1\}$, any score measure, any score $\gamma \in (\frac{1}{2}, 1)$, and every $n \geq 1$,*

$$R_{\gamma^n}(f^n) = \Omega\left(n \cdot \overline{\overline{R}}_{\gamma}(f)\right).$$

Proof. Combining Theorem 7 with our earlier results relating the different complexity measures, we obtain

$$\begin{aligned} R_{\gamma^n}(f^n) &\geq \overline{\overline{sR}}_{\gamma^n}(f^n) && \text{(Proposition 11)} \\ &\geq \Omega\left(n \overline{\overline{sR}}_{\gamma}(f)\right) && \text{(Theorem 7)} \\ &\geq \Omega\left(n \left(\frac{1}{2} \overline{\overline{R}}_{\gamma}(f)\right)\right) && \text{(Proposition 10)} \\ &= \Omega\left(n \overline{\overline{R}}_{\gamma}(f)\right). \end{aligned}$$

□

3.2 Proof of the Upper Bound

As a first step in the proof of the upper bound of Theorem 1, we establish a basic result showing that it is possible to bound both the average and worst-case cost of algorithms effectively.

Proposition 13. *Fix any $\delta > 0$. For any γ in the range $\frac{1}{2} + \delta \leq \gamma \leq 1 - \delta$, $f: \{0, 1\}^n \rightarrow \{0, 1\}$, and distribution μ , there exists an algorithm A that simultaneously satisfies*

$$\begin{aligned} \overline{\text{success}}_{\mu}^f(A) &\geq \gamma + \frac{\delta}{6}, \\ \overline{\text{cost}}^{\mu}(A) &\leq 4 \cdot \overline{\overline{R}}_{\gamma}(f), \quad \text{and} \\ \text{cost}(A) &\leq \frac{24}{\delta} \cdot \overline{\overline{R}}_{\gamma}(f). \end{aligned}$$

Proof. Let A_0 be an algorithm which satisfies $\overline{\text{success}}_{\mu}^f(A_0) \geq \gamma$ and $\overline{\text{cost}}^{\mu}(A_0) \leq \overline{\overline{R}}_{\gamma}^{\mu}(f) \leq \overline{\overline{R}}_{\gamma}(f)$. By the success amplification Lemma 8, there exists an algorithm A_1 with $\overline{\text{success}}_{\mu}^f(A_1) \geq \gamma + \frac{\delta}{3}$ and $\overline{\text{cost}}^{\mu}(A_1) \leq 4 \cdot \overline{\text{cost}}^{\mu}(A_0) \leq 4 \cdot \overline{\overline{R}}_{\gamma}(f)$.

Define A to be the algorithm that simulates A_1 but aborts and guesses the value of f on the input whenever A_1 's cost exceeds $\frac{6}{\delta} \cdot \overline{\text{cost}}^{\mu}(A_1)$. By construction, the average- and worst-case costs of A satisfy the conclusion of the proposition. And by Markov's inequality, the probability that A aborts A_1 's execution is at most $\frac{\delta}{6}$, so its success probability is at least $(\gamma + \frac{\delta}{3}) - \frac{\delta}{6} = \gamma + \frac{\delta}{6}$. □

We also use the following variant of the Azuma–Hoeffding inequality.

Lemma 14. *Let (X_1, \dots, X_n) be a sequence of random variables over \mathbb{R}^n that satisfies $\mathbb{E}[X_k \mid X_1, \dots, X_{k-1}] \leq \alpha$ and $|X_k| \leq c$ for all $k = 1, \dots, n$. Then*

$$\Pr\left[\sum_{i=1}^n X_i > \alpha n + t\right] \leq e^{-t^2/2c^2n}.$$

Proof. The sequence (Y_0, Y_1, \dots, Y_n) with $Y_0 = 0$ and $Y_k = \sum_{i=1}^k (X_i - \alpha)$ forms a submartingale with bounded differences $|Y_{k+1} - Y_k| \leq c$, and so by the standard form of the Azuma–Hoeffding inequality,

$$\Pr\left[\sum_{i=1}^n X_i > \alpha n + t\right] = \Pr[Y_n > t] \leq e^{-t^2/2c^2n}. \quad \square$$

We are now ready to complete the proof of the upper bound for the Direct Product Theorem.

Theorem 15 (Upper Bound of [Theorem 1](#)). *For any function $f: \{0, 1\}^m \rightarrow \{0, 1\}$, success parameter γ in the range $\frac{1}{2} < \gamma < 1$, and $n \geq \frac{2(24)^2}{(\min\{2\gamma-1, 1-\gamma\})^2}$,*

$$R_{\gamma^n}(f^n) = O\left(n \cdot \overline{R}_{\gamma}(f)\right). \quad (7)$$

Proof. Fix $\delta = \min\{2\gamma - 1, 1 - \gamma\}$. By Yao's Minimax Principle, it suffices to show that for any fixed distribution μ , there is an algorithm with worst-case cost $O(n \overline{R}_{\gamma}(f))$ that computes f^n with score at least γ^n on inputs drawn from the distribution μ . Fix μ . Consider the algorithm B_0 that computes the value of f^n on inputs drawn from μ by simulating algorithms A_1, \dots, A_n on the n instances of the input in the following way.

First, let μ_1 denote the marginal of μ on the first instance of the input. By [Proposition 13](#), there exists an algorithm A_1 with $\overline{\text{success}}_f^{\mu_1}(A_1) \geq \gamma + \frac{\delta}{6}$, $\overline{\text{cost}}^{\mu_1}(A_1) \leq 4 \cdot \overline{R}_{\gamma}(f)$, and $\text{cost}(A_1) \leq \frac{24}{\delta} \overline{R}_{\gamma}(f)$. The algorithm B_0 executes A_1 on the first instance of the input. Then for $i = 2, 3, \dots, n$ considered in order, let μ_i denote the marginal of μ on the i th instance of the input conditioned on the bits of the first $i-1$ instances that have been revealed by the algorithms A_1, \dots, A_{i-1} that have been previously executed. There exists an algorithm A_i that satisfies $\overline{\text{success}}_f^{\mu_i}(A_i) \geq \gamma + \frac{\delta}{6}$, $\overline{\text{cost}}^{\mu_i}(A_i) \leq 4 \cdot \overline{R}_{\gamma}(f)$, and $\text{cost}(A_i) \leq \frac{24}{\delta} \overline{R}_{\gamma}(f)$. The algorithm B_0 executes A_i on the i th instance of the input.

By the guarantees on A_1, \dots, A_n , the score of B_0 is

$$\overline{\text{success}}_f^{\mu}(B_0) = \prod_{i=1}^n \overline{\text{success}}_f^{\mu_i}(A_i) \geq \left(\gamma + \frac{\delta}{6}\right)^n$$

and its expected cost satisfies

$$\overline{\text{cost}}^{\mu}(B_0) = \sum_{i=1}^n \overline{\text{cost}}^{\mu_i}(A_i) \leq 4n \cdot \overline{R}_{\gamma}(f).$$

To complete the proof of the upper bound in [Theorem 1](#), we must also bound the worst-case cost of our algorithm for f^n . Unfortunately, the cost of B_0 can be as large as $\frac{24}{\delta} n \cdot \overline{R}_{\gamma}(f)$, which is larger than our target bound by a factor of $O(\frac{1}{\delta})$.

So let us now define B to be the algorithm that simulates B_0 but aborts whenever the total cost of the simulation exceeds $9n \cdot \overline{R}_{\gamma}(f)$. The algorithm B outputs any value whenever it aborts. Let us write E to denote the event that B aborts, and let W denote the event that $B_0(x) = f^n(x)$. The probability that B succeeds on inputs drawn from μ is bounded below by

$$\begin{aligned} \Pr_{x \sim \mu} [B(x) = f^n(x)] &\geq \Pr_{x \sim \mu} [W \wedge \overline{E}] \\ &= \Pr_{x \sim \mu} [W] \cdot \Pr_{x \sim \mu} [\overline{E} \mid W]. \end{aligned}$$

We have already seen that $\Pr_{x \sim \mu} [W] = \overline{\text{success}}_f^{\mu}(B_0) \geq (\gamma + \frac{\delta}{6})^n$. It remains to bound the second term.

Define X_i to be a random variable that represents the cost of A_i when we run B_0 . The expected value of X_i conditioned on X_1, \dots, X_{i-1} and W satisfies

$$\begin{aligned} \mathbb{E}_{x \sim \mu} [X_i \mid X_1, \dots, X_{i-1}, W] &\leq \max_{\mu_i} \mathbb{E}_{x_i \sim \mu_i} [X_i \mid A_i(x_i) = f(x_i)] \\ &\leq \max_{\mu_i} \frac{\mathbb{E}_{x_i \sim \mu_i} [X_i]}{\Pr_{x_i \sim \mu_i} [A_i(x_i) = f(x_i)]} \\ &\leq \max_{\mu_i} \frac{\overline{\text{cost}}^{\mu_i}(A_i)}{\overline{\text{success}}_f^{\mu_i}(A_i)} \end{aligned}$$

Since the success probability score of any algorithm is bounded below by $\overline{\text{success}}_f^{\mu_i}(A_i) \geq \frac{1}{2}$, we conclude that

$$\mathbb{E}_{x \sim \mu} [X_i \mid X_1, \dots, X_{i-1}, W] \leq 2 \max_{\mu_i} \overline{\text{cost}}^{\mu_i}(A_i) \leq 8 \overline{\text{R}}_\gamma(f).$$

And each X_i is bounded by $0 \leq X_i \leq \frac{24}{\delta} \overline{\text{R}}_\gamma(f)$. By Lemma 14 with parameters $t = n \overline{\text{R}}_\gamma(f)$ and $c = \frac{24}{\delta} \overline{\text{R}}_\gamma(f)$, we then have

$$\Pr[E \mid W] \leq \Pr[X > \mathbb{E}[X \mid W] + n \cdot \overline{\text{R}}_\gamma(f) \mid W] \leq e^{-\frac{1}{2 \cdot (24)^2} \delta^2 n}.$$

When $n > 2(24)^2/\delta^2$, the last expression is smaller than e^{-1} . And in this case the overall success of B is at least $(\gamma + \frac{\delta}{6})^n e^{-1} \geq \gamma^n \cdot (1 + \frac{\delta}{6})^n e^{-1} \geq \gamma^n e^{\frac{\delta}{12} n - 1}$, which is at least γ^n when $n \geq 12/\delta$. \square

4 Proof of the Tensorization Lemma

Lemma 5 (Tensorization Lemma). *For any (possibly partial) function $f: \{0, 1\}^m \rightarrow \{0, 1\}$, any distribution μ on $\{0, 1\}^m$, and any parameter $n \geq 1$,*

$$\text{DS}_\alpha^{\mu^n}(f^n) = \text{DS}_\alpha^\mu(f)^n.$$

Proof. We start by proving that $\text{DS}_\alpha^{\mu^n}(f^n) \leq \text{DS}_\alpha^\mu(f)^n$. Let A be a randomized algorithm with discounted score $\text{ds}_{f,\alpha}^\mu(A) = \text{DS}_\alpha^\mu(f)$. Consider the algorithm B that computes f^n by running n independent instances of A sequentially on each of the n inputs. By the independence of the n inputs and of the n instances of A , the discounted score of B over μ^n is

$$\begin{aligned} \text{ds}_{f^n,\alpha}^{\mu^n}(B) &= \mathbb{E}_{\substack{x_1, \dots, x_n \sim \mu \\ D^{(1)}, \dots, D^{(n)} \sim A}} \left[\prod_{i=1}^n \text{score}_f^\mu(D^{(i)}(x_i)) \cdot e^{-\alpha \text{cost}(D^{(i)}(x_i))} \right] \\ &= \prod_{i=1}^n \mathbb{E}_{\substack{x_i \sim \mu \\ D^{(i)} \sim A}} \left[\text{score}_f^\mu(D^{(i)}(x_i)) \cdot e^{-\alpha \text{cost}(D^{(i)}(x_i))} \right] \\ &= \text{ds}_{f,\alpha}^\mu(A)^n. \end{aligned}$$

Therefore, $\text{DS}_\alpha^{\mu^n}(f^n) \leq \text{ds}_{f^n,\alpha}^{\mu^n}(B) = \text{ds}_{f,\alpha}^\mu(A)^n = \text{DS}_\alpha^\mu(f)^n$, as we wanted to show.

We now prove that $\text{DS}_\alpha^{\mu^n}(f^n) \geq \text{DS}_\alpha^\mu(f)^n$. Let B be an algorithm that computes f^n and has discounted score $\text{ds}_{f^n,\alpha}^{\mu^n}(B) = \text{DS}_\alpha^{\mu^n}(f^n)$. Without loss of generality, we can assume that B is deterministic. We define a randomized algorithm A for f in the following way.

The algorithm A is defined with respect to a specific distribution ν on the domain of f^n that depends on f , μ , D , and α in a way that will be specified later. The behaviour of A is defined as follows. First, it picks $i \in [n]$ uniformly at random. Next, it runs the algorithm A_i which works by running B on $n-1$ fake inputs and the real input to f , with the real input placed in position i . When B makes a query inside the i th input to f , the algorithm A_i queries the real input. And when B makes a query to the j th input for some $j \neq i$, the answer to the query is generated artificially by the algorithm A_i ; it generates the query answer from the conditional distribution ν conditioned on the transcript of B so far. (That is, it samples a string y from ν conditioned on the string being consistent with all the queries answered so far, and it determines the value of the query by making the corresponding query to y .)

Let π_i be the distribution over the leaves of D reached by this simulation by A_i , assuming the real input (the one at position i) was sampled from μ . Extend π_i to a distribution over the domain

of f^n by replacing each leaf of B with a sample from μ^n conditioned on reaching that leaf. The behaviour of A_i on an input sampled from μ is then the same as the behaviour of B on an input sampled from π_i . For an algorithm acting on n inputs to f , let cost_i denote the number of queries to the i th input. We now write

$$\text{ds}_{f,\alpha}^\mu(A) = \mathbb{E}_{i \sim [n]} \left[\sum_{\ell \in B} \pi_i(\ell) \cdot \text{score}_f^\mu(\ell_i) e^{-\alpha \text{cost}_i(\ell)} \right] = \sum_{\ell \in B} \mathbb{E}_{i \sim [n]} \left[\pi_i(\ell) \cdot \text{score}_f^\mu(\ell_i) e^{-\alpha \text{cost}_i(\ell)} \right].$$

The positivity of the score measures guarantees that each term of the sum is non-negative. By applying the AM-GM inequality to each term of the sum, the last expression is bounded below by

$$\sum_{\ell \in B} \mathbb{E}_{i \sim [n]} \left[\pi_i(\ell) \cdot \text{score}_f^\mu(\ell_i) e^{-\alpha \text{cost}_i(\ell)} \right] \geq \sum_{\ell} \left(\prod_{i=1}^n \pi_i(\ell) \cdot \text{score}_f^\mu(\ell_i) e^{-\alpha \text{cost}_i(\ell)} \right)^{1/n}.$$

We next observe that $\sum_i \text{cost}_i(\ell) = \text{cost}(\ell)$. By definition, we also have $\prod_i \text{score}_f^\mu(\ell_i) = \text{score}_{f^n}^{\mu^n}(\ell)$. And then the key observation at the heart of the proof is that we also have

$$\prod_{i=1}^n \pi_i(\ell) = \nu(\ell)^{n-1} \mu^n(\ell).$$

To see why this last identity holds, note that we can write $\pi_i(\ell) = \prod_{(v,w)} \pi_i(w|v)$ where the product is over all the directed edges (v,w) from the root of the tree to the leaf ℓ . By our construction,

$$\pi_i(w|v) = \begin{cases} \mu^n(w|v) & \text{if node } v \text{ queries the } i\text{th instance} \\ \nu(w|v) & \text{otherwise.} \end{cases}$$

So we have

$$\begin{aligned} \prod_{i=1}^n \pi_i(\ell) &= \prod_{(v,w)} \prod_i \pi_i(w|v) = \prod_{(v,w)} \mu^n(w|v) \nu(w|v)^{n-1} \\ &= \prod_{(v,w)} \mu^n(w|v) \left(\prod_{(v,w)} \nu(w|v) \right)^{n-1} = \mu^n(\ell) \nu(\ell)^{n-1}. \end{aligned}$$

Combining all of the above, we then end up with

$$\text{ds}_{f,\alpha}^\mu(A) \geq \sum_{\ell} \left(\nu(\ell)^{n-1} \mu^n(\ell) \cdot \text{score}_{f^n}^{\mu^n}(\ell) e^{-\alpha \text{cost}(\ell)} \right)^{1/n}.$$

Define ν by setting

$$\nu(y) = \frac{\mu^n(y) \cdot \text{score}_{f^n}^{\mu^n}(B(y)) e^{-\alpha \text{cost}(B(y))}}{\sum_z \mu^n(z) \cdot \text{score}_{f^n}^{\mu^n}(B(z)) e^{-\alpha \text{cost}(B(z))}}.$$

This gives the corresponding distribution on leaves defined by

$$\nu(\ell) = \frac{\mu^n(\ell) \cdot \text{score}_{f^n}^{\mu^n}(\ell) e^{-\alpha \text{cost}(\ell)}}{\sum_{\ell'} \mu^n(\ell') \cdot \text{score}_{f^n}^{\mu^n}(\ell') e^{-\alpha \text{cost}(\ell')}},$$

and so we obtain

$$\text{ds}_{f,\alpha}^\mu(A) \geq \left(\sum_{\ell} \mu^n(\ell) \cdot \text{score}_{f^n}^{\mu^n}(\ell) e^{-\alpha \text{cost}(\ell)} \right)^{1/n} = \text{ds}_{f^n,\alpha}^{\mu^n}(B)^{1/n}. \quad \square$$

5 Proof of the Equivalence Lemma

5.1 Proof Overview

Both the upper and lower bounds in [Lemma 6](#) rely on an alternative characterization of discounted score and score-weighted cost in terms of a reweighted distribution on leaves. Let us first introduce this distribution.

For any randomized algorithm R , let us denote the distribution π_R on the leaves of the deterministic decision trees in the support of R obtained by taking the probability that we reach this leaf when $D \sim R$ and $x \sim \mu$ and rescaling it by the weight $\overline{\text{score}}_f^\mu(\ell)$. Namely, letting $\mu(x)$ and $R(D)$ denote the probability that $x \sim \mu$ and $D \sim R$, respectively, we define the rescaled probability of the leaf ℓ of D to be

$$\pi_R(\ell) = \frac{\mu(\ell)R(D) \cdot \overline{\text{score}}_f^\mu(\ell)}{\overline{\text{score}}_f^\mu(R)}.$$

The discounted score of the randomized algorithm R satisfies

$$\text{ds}_{f,\alpha}^\mu(R) = \overline{\text{score}}_f^\mu(R) \cdot \mathbb{E}_{\ell \sim \pi_R} [e^{-\alpha \text{cost}(\ell)}]$$

and the success-conditioned cost of an algorithm satisfies

$$\overline{\text{scost}}_f^\mu(R) = \mathbb{E}_{\ell \sim \pi_R} [\text{cost}(\ell)].$$

This means that we can express the maximum discounted score of f as

$$\text{DS}_\alpha^\mu(f) = \max_R \overline{\text{score}}_f^\mu(R) \cdot \mathbb{E}_{\ell \sim \pi_R} [e^{-\alpha \text{cost}(\ell)}]$$

and its distributional success-conditioned query complexity of f can be expressed as

$$\overline{\text{SR}}_\gamma^\mu(f) = \min_{R: \overline{\text{score}}_f^\mu(R) \geq \gamma} \mathbb{E}_{\ell \sim \pi_R} [\text{cost}(\ell)].$$

With this characterization, the lower bound in [Lemma 6](#) is an immediate consequence of Jensen's inequality.

The upper bound requires a few more ingredients. The first is [Proposition 28](#), showing that for every success parameter γ , there exists a discount parameter α^* for which the maximum discounted score $\text{DS}_{\alpha^*}^\mu(f)$ is achieved by a randomized algorithm R with average score $\overline{\text{score}}_f^\mu(R)$ exactly equal to γ . Choosing the parameter α^* and corresponding algorithm R^* in the conclusion of the proposition, we obtain the identity

$$\text{DS}_{\alpha^*}^\mu(f) = \gamma \cdot \mathbb{E}_{\pi_{R^*}} [e^{-\alpha^* \text{cost}(\ell)}].$$

What we would like to use at this point is a “reverse Jensen inequality” to obtain a bound of the form

$$\mathbb{E}_{\pi_{R^*}} [e^{-\alpha^* \text{cost}(\ell)}] \stackrel{?}{\leq} e^{-c \alpha^* \cdot \mathbb{E}_{\pi_{R^*}} [\text{cost}(\ell)]}$$

for some constant c . No such inequality holds in general, but we do have the fact that when y is a random variable bounded by some interval $[0, d]$, then convexity of the e^{-y} function implies that $e^{-y} \leq (1 - \frac{y}{d})e^{-0} + \frac{y}{d}e^{-d}$ and so

$$\mathbb{E}[e^{-y}] \leq 1 - \frac{1 - e^{-d}}{d} \mathbb{E}[y] \leq e^{-\frac{1 - e^{-d}}{d} \mathbb{E}[y]}.$$

Applying this inequality to the random variable $y = \alpha^* \text{cost}(R^*(x))$ then gives exactly the type of inequality we are aiming for if we have the guarantee that the cost of the leaves in R^* is bounded above by $O(1/\alpha^*)$.

In general, however, we do not have any such guarantee on the maximum cost of the algorithm R^* . We can get around this limitation by considering a truncation of R^* that yields a different algorithm R^\dagger . The key to completing the argument with this approach is to show that such a truncation does not affect the average cost by more than a constant multiplicative factor. And this is where the fact that R^* has maximum discounted score for the parameter α^* comes into play: we can argue that any such algorithm cannot have much of its average cost (even in the rescaled distribution π_{R^*}) on high-cost leaves. The proof of this last fact is obtained by a contradiction argument, showing that if the last condition holds, we can define an algorithm with strictly better discounted score by truncating part of the tree of R^* , contradicting the fact that it has maximal discounted score.

5.2 The Proof

Lemma 6 (Equivalence Lemma). *For every (possibly partial) function $f: \{0,1\}^m \rightarrow \{0,1\}$, every distribution μ over the domain of f , every success parameter γ , and every discount factor $\alpha \geq 0$,*

$$\overline{\text{SR}}_\gamma^\mu(f) \geq \frac{1}{\alpha} \log \frac{\text{DS}_\alpha^\mu(f)}{\gamma}.$$

Moreover, this lower bound is best possible: for every f , μ , and γ , there exists a discount factor α^* for which

$$\overline{\text{SR}}_\gamma^\mu(f) = O\left(\frac{1}{\alpha^*} \log \frac{\text{DS}_{\alpha^*}^\mu(f)}{\gamma}\right).$$

Proof. By the characterization of maximum discounted score introduced in the proof overview,

$$\begin{aligned} \text{DS}_\alpha^\mu(f) &= \max_R \overline{\text{score}}_f^\mu(R) \cdot \mathbb{E}_{\ell \sim \pi_R} [e^{-\alpha \text{cost}(\ell)}] \\ &\geq \max_{R: \overline{\text{score}}_f^\mu(R) \geq \gamma} \overline{\text{score}}_f^\mu(R) \cdot \mathbb{E}_{\pi_R} [e^{-\alpha \text{cost}(\ell)}] \\ &\geq \gamma \cdot \max_{R: \overline{\text{score}}_f^\mu(R) \geq \gamma} \mathbb{E}_{\pi_R} [e^{-\alpha \text{cost}(\ell)}]. \end{aligned}$$

By Jensen's inequality and the convexity of the function $y \mapsto e^{-\alpha y}$, for every algorithm R we have the inequality $\mathbb{E}_{\pi_R} [e^{-\alpha \text{cost}(\ell)}] \geq e^{-\alpha \mathbb{E}_{\pi_R} [\text{cost}(\ell)]}$. Therefore,

$$\begin{aligned} \text{DS}_\alpha^\mu(f) &\geq \gamma \cdot \max_{R: \overline{\text{score}}_f^\mu(R) \geq \gamma} e^{-\alpha \mathbb{E}_{\pi_R} [\text{cost}(\ell)]} \\ &= \gamma \cdot e^{-\alpha \cdot \min_{R: \overline{\text{score}}_f^\mu(R) \geq \gamma} \mathbb{E}_{\pi_R} [\text{cost}(\ell)]} \\ &= \gamma \cdot e^{-\alpha \overline{\text{SR}}_\gamma^\mu(f)}. \end{aligned}$$

This completes the proof of the lower bound in the lemma.

We now complete the proof of the upper bound. Let α^* be the discount rate and R^* be the associated algorithm in the conclusion of [Proposition 28](#). These objects satisfy

$$\text{DS}_{\alpha^*}^\mu(f) = \overline{\text{score}}_f^\mu(R^*) \cdot \mathbb{E}_{\ell \sim \pi_{R^*}} [e^{-\alpha^* \text{cost}(\ell)}] = \gamma \cdot \mathbb{E}_{\ell \sim \pi_{R^*}} [e^{-\alpha^* \text{cost}(\ell)}].$$

Let d be a depth parameter whose value will be set later. The trivial inequality $\min\{\text{cost}(\ell), d\} \leq \text{cost}(\ell)$ implies that $\text{DS}_{\alpha^*}^\mu(f)$ can be bounded above by

$$\text{DS}_{\alpha^*}^\mu(f) \leq \gamma \cdot \mathbb{E}_{\ell \sim \pi_{R^*}} [e^{-\alpha^* \min\{\text{cost}(\ell), d\}}].$$

By convexity of the function $e^{-\alpha^* y}$, when y is bounded by $0 \leq y \leq d$, then

$$e^{-\alpha^* y} \leq \left(1 - \frac{y}{d}\right) \cdot e^{-\alpha^* \cdot 0} + \frac{y}{d} \cdot e^{-\alpha^* d} = 1 - \frac{1 - e^{-\alpha^* d}}{d} \cdot y.$$

Applying this inequality to the random variable $y = \min\{\text{cost}(\ell), d\}$, we obtain

$$\mathbb{E}_{\ell \sim \pi_{R^*}} [e^{-\alpha^* \min\{\text{cost}(\ell), d\}}] \leq 1 - \frac{1 - e^{-\alpha^* d}}{d} \cdot \mathbb{E}_{\pi_{R^*}} [\min\{\text{cost}(\ell), d\}] \leq e^{-\frac{1 - e^{-\alpha^* d}}{d} \cdot \mathbb{E}_{\pi_{R^*}} [\min\{\text{cost}(\ell), d\}]}.$$

Setting $d = 2 \ln(2)/\alpha^*$, we observe that

$$\frac{1 - e^{-\alpha^* d}}{d} = \frac{1 - 2^{-2}}{2 \ln(2)/\alpha^*} = \frac{3}{8 \ln(2)} \cdot \alpha^*.$$

Combining all of the above gives us the inequality

$$\text{DS}_{\alpha^*}^\mu(f) \leq \gamma \cdot e^{-c' \alpha^* \mathbb{E}_{\pi_{R^*}} [\text{cost}(\lfloor D \rfloor_d(x))]}$$

for the constant $c' = 3/8 \ln(2)$. To complete the proof of the upper bound, it remains to bound $\mathbb{E}_{\pi_{R^*}} [\min\{\text{cost}(\ell), d\}]$ from below by a multiplicative constant times $\mathbb{E}_{\pi_{R^*}} [\text{cost}(\ell)]$.

Claim 16. *With R^* and d as defined above, $\mathbb{E}_{\pi_{R^*}} [\min\{\text{cost}(\ell), d\}] \geq \frac{1}{3} \mathbb{E}_{\pi_{R^*}} [\text{cost}(\ell)]$.*

For any depth τ , let us write $X_{>\tau}$ to denote the event that $\text{cost}(\ell) > \tau$ and, in a slight abuse of notation, let us also write $X_{>\tau}$ to denote the indicator random variable corresponding to this event. Similarly, let $X_{\leq\tau}$ denote the event that $\text{cost}(\ell) \leq \tau$ and its indicator random variable.

Let us write $L(\mu, R^*)$ to denote the distribution on the leaf $D(x)$ reached when we draw $x \sim \mu$ and $D \sim R^*$. We can write the discounted score of R^* as

$$\begin{aligned} \text{ds}_{f, \alpha^*}^\mu(R^*) &= \mathbb{E}_{\ell \sim L(\mu, R^*)} [\overline{\text{score}}_f^\mu(\ell) \cdot e^{-\alpha^* \text{cost}(\ell)} X_{\leq\tau}] + \mathbb{E}_{\ell \sim L(\mu, R^*)} [\overline{\text{score}}_f^\mu(\ell) \cdot e^{-\alpha^* \text{cost}(\ell)} X_{>\tau}] \\ &= \mathbb{E}_{\ell \sim L(\mu, R^*)} [\overline{\text{score}}_f^\mu(\ell) \cdot e^{-\alpha^* \text{cost}(\ell)} X_{\leq\tau}] + \gamma \mathbb{E}_{\pi_{R^*}} [e^{-\alpha^* \text{cost}(\ell)} X_{>\tau}]. \end{aligned}$$

Let $\lfloor R^* \rfloor_\tau$ be the randomized algorithm by truncating R^* at depth τ . The truncation is completed by replacing each internal node at depth τ in a tree D in the support of R^* with a leaf to obtain the tree $\lfloor D \rfloor_\tau$. For every leaf ℓ in a decision tree D in the support of R^* , let us write $\lfloor \ell \rfloor_\tau$ to denote the leaf in $\lfloor D \rfloor_\tau$ reached by the inputs that lead to ℓ in D . The cost of these leaves satisfies $\text{cost}(\lfloor \ell \rfloor_\tau) = \min\{\text{cost}(\ell), \tau\}$. And the discounted score of the algorithm $\lfloor R^* \rfloor_\tau$ is bounded below by

$$\text{ds}_{f, \alpha^*}^\mu(\lfloor R^* \rfloor_\tau) \geq \mathbb{E}_{\ell \sim L(\mu, R^*)} [\overline{\text{score}}_f^\mu(\ell) e^{-\alpha^* \text{cost}(\ell)} X_{\leq\tau}] + \frac{1}{2} e^{-\alpha^* \tau} \Pr_{\ell \sim L(\mu, R^*)} [X_{>\tau}]$$

where the second term uses the bounded range property of scoring measures which guarantees that the score of every leaf (including in particular the ones introduced during the truncation process) is at least $\frac{1}{2}$. Since R^* has maximum discounted score, we must have $\text{ds}_{f, \alpha^*}^\mu(R^*) \geq \text{ds}_{f, \alpha^*}^\mu(\lfloor R^* \rfloor_\tau)$. And so the above decompositions of the discounted scores of R^* and $\lfloor R^* \rfloor_\tau$ imply that

$$\gamma \mathbb{E}_{\ell \sim \pi_{R^*}} [e^{-\alpha^* \text{cost}(\ell)} X_{>\tau}] \geq \frac{1}{2} e^{-\alpha^* \tau} \Pr_{\ell \sim L(\mu, R^*)} [X_{>\tau}].$$

The right-hand side of this inequality can be further bounded below by noting that

$$\Pr_{\ell \sim L(\mu, R^*)} [X_{>\tau}] \geq \mathbb{E}_{\ell \sim L(\mu, R^*)} [\text{score}_f^\mu(\ell) X_{>\tau}] = \text{score}_f^\mu(R^*) \mathbb{E}_{\ell \sim \pi_{R^*}} [X_{>\tau}] = \gamma \Pr_{\ell \sim \pi_{R^*}} [X_{>\tau}].$$

And the left-hand side can be bounded above by using

$$\mathbb{E}_{\ell \sim \pi_{R^*}} [e^{-\alpha^* \text{cost}(\ell)} X_{>\tau}] \leq \left(\Pr_{\pi_{R^*}} [X_{>\tau}] - \Pr_{\pi_{R^*}} [X_{>\tau+d}] \right) e^{-\alpha^* \tau} + \Pr_{\pi_{R^*}} [X_{>\tau+d}] e^{-\alpha^* (\tau+d)}.$$

Combining the last three inequalities, we obtain

$$\Pr_{\pi_{R^*}} [X_{>\tau+d}] (1 - e^{-\alpha^* d}) \leq \frac{1}{2} \Pr_{\pi_{R^*}} [X_{>\tau}].$$

Recalling that $d = 2 \ln(2)/\alpha^*$, this gives $\frac{3}{4} \Pr_{\pi_{R^*}} [X_{>\tau+d}] \leq \frac{1}{2} \Pr_{\pi_{R^*}} [X_{>\tau}]$ or, equivalently,

$$\Pr_{\pi_{R^*}} [X_{>\tau+d}] \leq \frac{2}{3} \Pr_{\pi_{R^*}} [X_{>\tau}].$$

We now use this inequality to complete the proof of the claim.

The expected cost of R^* under the distribution π_{R^*} can be expressed as

$$\mathbb{E}_{\pi_{R^*}} [\text{cost}(\ell)] = \mathbb{E}_{\pi_{R^*}} [\text{cost}(\ell) X_{\leq d}] + \mathbb{E}_{\pi_{R^*}} [\text{cost}(\ell) X_{>d}].$$

The latter term can be bounded above by

$$\begin{aligned} \mathbb{E}_{\pi_{R^*}} [\text{cost}(\ell) X_{>d}] &\leq \sum_{k \geq 1} (k+1)d \left(\Pr_{\pi_{R^*}} [X_{>kd}] - \Pr_{\pi_{R^*}} [X_{>(k+1)d}] \right) \\ &= d \sum_{k \geq 1} \Pr_{\pi_{R^*}} [X_{>kd}]. \end{aligned}$$

From the inequality in the last paragraph, for every $k \geq 2$,

$$\Pr_{\pi_{R^*}} [X_{>kd}] \leq \frac{2}{3} \Pr_{\pi_{R^*}} [X_{>(k-1)d}] \leq \left(\frac{2}{3}\right)^{k-1} \Pr_{\pi_{R^*}} [X_{>d}].$$

So $\mathbb{E}_{\pi_{R^*}} [\text{cost}(\ell) X_{>d}] \leq d \sum_{k \geq 0} \left(\frac{2}{3}\right)^k \Pr_{\pi_{R^*}} [X_{>d}]$ and

$$\mathbb{E}_{\pi_{R^*}} [\text{cost}(\ell)] \leq \mathbb{E}_{\pi_{R^*}} [\text{cost}(\ell) X_{\leq d}] + 3d \Pr_{\pi_{R^*}} [X_{>d}].$$

Meanwhile, the expected cost of $\lfloor R^* \rfloor_d$ under π_{R^*} can be expressed as

$$\mathbb{E}_{\pi_{R^*}} [\text{cost}(\lfloor \ell \rfloor_d)] = \mathbb{E}_{\pi_{R^*}} [\text{cost}(\ell) X_{\leq d}] + d \Pr_{\pi_{R^*}} [X_{>d}].$$

Therefore, $\mathbb{E}_{\pi_{R^*}} [\text{cost}(\ell)] \leq 3 \mathbb{E}_{\pi_{R^*}} [\text{cost}(\lfloor \ell \rfloor_d)] = 3 \mathbb{E}_{\pi_{R^*}} [\min\{\text{cost}(\ell), d\}]$, as we wanted to show. \square

6 List-Decoding Direct Product Theorems

6.1 Preliminary Lemmas

Before we can complete the proof of [Theorem 2](#), we need to establish a few information-theoretic facts. The *entropy* of the random variable $X \sim \pi$ over a finite domain is

$$H(X) = \sum_x \pi(x) \log \frac{1}{\pi(x)} = \mathbb{E}_{X \sim \pi} [\log \frac{1}{\pi(X)}].$$

The *cross-entropy* of $X \sim \pi$ with respect to $Z \sim \sigma$ over the same domain is

$$H(Z, X) = \mathbb{E}_{X \sim \sigma} [\log \frac{1}{\pi(X)}].$$

We start with a basic inequality on the entropy of collections of independent Bernoulli random variables.

Lemma 17. *Let X_1, \dots, X_n and Z_1, \dots, Z_n be $2n$ independent Bernoulli random variables with $X_i \sim \text{Ber}(\frac{1+\delta_i}{2})$ and $Z_i \sim \text{Ber}(\frac{1+\tau_i}{2})$ for $i = 1, \dots, n$. If $X = (X_1, \dots, X_n)$ and $Z = (Z_1, \dots, Z_n)$ satisfy $H(Z) \geq n - \alpha$ and $H(Z, X) \leq n - \beta$ for some $\beta \leq \alpha$, then*

$$H(X) \leq n - \frac{\beta^2}{16\alpha}.$$

Proof. If there is any index i for which $\delta_i > \tau_i$, then the random variable X' drawn as X except with bias τ_i for X_i also satisfies the two constraints since $H(Z_i, X_i) \geq H(Z_i) = H(Z_i, X'_i)$ and it has strictly larger entropy $H(X') > H(X)$. Therefore, without loss of generality, we can consider only the case where $\delta_i \leq \tau_i$ for each $i = 1, \dots, n$.

The entropy of X_i satisfies

$$1 - H(X_i) = 1 - h\left(\frac{1+\delta_i}{2}\right) = \frac{1}{2\ln 2} \left(\delta_i^2 + \frac{\delta_i^4}{6} + \frac{\delta_i^6}{15} + \dots \right).$$

So $\frac{\delta_i^2}{2\ln 2} \leq 1 - H(X_i) \leq \delta_i^2$ and by adding the entropy of each coordinate of X ,

$$\frac{\|\delta\|_2^2}{2\ln 2} \leq n - H(X) \leq \|\delta\|_2^2.$$

Similarly, $n - H(Z) \geq \|\tau\|_2^2/2\ln 2$. With the condition in the lemma statement, we therefore have the bound

$$\|\tau\|_2^2 \leq 2\ln 2 \cdot \alpha.$$

We also have

$$1 - H(Z_i, X_i) = 1 - H(X_i) + \frac{\tau_i - \delta_i}{2} \log\left(\frac{1+\delta_i}{1-\delta_i}\right).$$

Since $\log(1+x) \leq x/\ln(2)$, the last term in this expression is bounded by

$$\frac{\tau_i - \delta_i}{2} \log\left(\frac{1+\delta_i}{1-\delta_i}\right) = \frac{\tau_i - \delta_i}{2} \log\left(1 + \frac{2\delta_i}{1-\delta_i}\right) \leq \frac{\tau_i - \delta_i}{2} \cdot \frac{2\delta_i}{\ln(2)(1-\delta_i)} = \frac{\tau_i - \delta_i}{1-\delta_i} \cdot \frac{\delta_i}{\ln 2} \leq \frac{\tau_i \delta_i}{\ln 2},$$

and we obtain

$$n - H(Z, X) \leq n - H(X) + \frac{\langle \tau, \delta \rangle}{\ln 2} \leq \|\delta\|_2^2 + \frac{\|\tau\|_2 \|\delta\|_2}{\ln 2} \leq \frac{2}{\ln 2} \|\tau\|_2 \|\delta\|_2$$

where the last inequality uses the fact that $\|\delta\|_2 \leq \|\tau\|_2 \leq \frac{\|\tau\|_2}{\ln 2}$.

The above inequalities and the conditions of the lemma now give

$$\beta^2 \leq (n - H(Z, X))^2 \leq \frac{4\|\tau\|_2^2 \|\delta\|_2^2}{\ln^2 2} \leq \frac{8}{\ln 2} \alpha \|\delta\|_2^2.$$

These inequalities imply that $\frac{\|\delta\|_2^2}{\ln 2} \geq \frac{\beta^2}{8\alpha}$. Therefore,

$$n - H(X) \geq \frac{\|\delta\|_2^2}{2\ln 2} \geq \frac{\beta^2}{16\alpha}.$$

□

We use the above bound to obtain the following result.

Lemma 18. *Fix any $n \in \mathbb{N}$, $\ell \in \{1, \dots, n\}$, and $\epsilon \in [0, \frac{1}{4}]$. If p_1, \dots, p_n are biases and $L \subseteq \{0, 1\}^n$ is a list of $2^{n-\ell}$ elements such that when $X_i \sim \text{Ber}(p_i)$ are drawn independently, $X = (X_1, \dots, X_n)$ satisfies*

$$\Pr[X \in L] \geq 2^{-\epsilon\ell},$$

then $H(X) \leq n - \frac{(1-\epsilon)^2\ell}{320}$.

Proof. Fix $\alpha = 2^{-\frac{1-\epsilon}{2}}$ and $d = \alpha 2^{(1-\epsilon)\ell-n}$. Consider the disjoint subsets B_0, B_1, B_2, \dots of L where for each $i \geq 0$,

$$B_i = \{x \in L : 2^i d < p(x) \leq 2^{i+1} d\}.$$

By the hypothesis of the lemma, $\sum_{x \in L} p(x) \geq 2^{-\epsilon\ell}$. And by construction, the only elements in L that are not in one of the bins B_i , $i \geq 0$ are the ones for which $p(x) \leq d = \alpha 2^{(1-\epsilon)\ell-n}$, so the total mass of the excluded elements is at most $2^{n-\ell} \cdot \alpha 2^{(1-\epsilon)\ell-n} = \alpha 2^{-\epsilon\ell}$. This means that the mass of the points in the union of the B_i sets is bounded below by

$$\sum_{i \geq 0} \sum_{x \in B_i} p(x) \geq (1 - \alpha) 2^{-\epsilon\ell}.$$

This lower bound implies that there must be an index $i \geq 0$ for which the bin B_i has total mass

$$\sum_{x \in B_i} p(x) \geq (1 - \alpha) 2^{-\epsilon\ell} \cdot 2^{-(i+1)} \quad (8)$$

otherwise we would have $\sum_{i \geq 0} \sum_{x \in B_i} p(x) < (1 - \alpha) 2^{-\epsilon\ell} \cdot 2^{-(i+1)} = (1 - \alpha) 2^{-\epsilon\ell}$, contradicting the lower bound on the mass of points in the union of the bins. Fix an index i that satisfies (8) and let D be the uniform distribution on the bin B_i .

The upper bound on the mass of each element in B_i and the lower bound on the total mass of this bin imply that the number of points in B_i is bounded below by

$$|B_i| \geq \frac{(1 - \alpha) 2^{-\epsilon\ell} \cdot 2^{-(i+1)}}{2^{i+1} d} = \frac{1 - \alpha}{\alpha} \cdot 2^{n-\ell} \cdot 2^{-2(i+1)}.$$

So the entropy of D is bounded below by

$$H(D) = \log |B_i| \geq n - (\ell + 2(i+1) - \log \frac{1-\alpha}{\alpha}).$$

When $\epsilon \leq \frac{1}{4}$, then $\log \frac{1-\alpha}{\alpha} = \log(2^{-\frac{1-\epsilon}{2}} - 1) \geq -2$, so in this case the lower bound simplifies to $H(D) \geq n - (\ell + 4 + 2i)$. Meanwhile, the lower bound on the mass of the elements in B_i implies that

$$H(D, X) = -\mathbb{E}_{x \sim D}[\log p(x)] \leq -\log(2^i d) = -\log(\alpha 2^{(1-\epsilon)\ell-n} 2^i) = n - ((1 - \epsilon)\ell - \log \frac{1}{\alpha} + i).$$

Since $\log \frac{1}{\alpha} = \frac{1-\epsilon}{2}$, we can simplify this bound to $H(D, X) \leq n - ((1 - \epsilon)\ell - \frac{1-\epsilon}{2} + i)$.

Consider now the Bernoulli random variables $Z_i \sim \text{Ber}(q_i)$ with parameters $q_i = \Pr_{x \in D}[x_i = 1]$ for $i = 1, \dots, n$. Each Z_i is equal to the marginal D_i of D on the i th coordinate. The random variable Z satisfies

$$H(Z) = \sum_{i=1}^n H(Z_i) = \sum_{i=1}^n H(D_i) \geq \sum_{i=1}^n H(D_i | D_{<i}) = H(D)$$

and, since p is a product distribution,

$$H(Z, X) = - \sum_{i=1}^n \mathbb{E}_{x_i \sim Z_i} [\log p_i(x_i)] = - \sum_{i=1}^n \mathbb{E}_{x_i \sim D_i} [\log p_i(x_i)] = H(D, X).$$

Applying [Lemma 17](#) to X and Z , we obtain

$$H(X) \leq n - \frac{((1-\epsilon)\ell - \frac{1-\epsilon}{2} + i)^2}{16(\ell + 4 + 2i)}.$$

Since the ratio $\frac{((1-\epsilon)\ell - \frac{1-\epsilon}{2} + i)^2}{16(\ell + 4 + 2i)}$ is monotone in i when $\ell \geq 1$ and $i \geq 0$, the entropy of X is bounded by

$$H(X) \leq n - \frac{((1-\epsilon)\ell - \frac{1-\epsilon}{2})^2}{16(\ell + 4)} = n - \frac{(1-\epsilon)^2}{16} \cdot \frac{(\ell - \frac{1}{2})^2}{\ell + 4}.$$

Note that when $\ell \geq 1$, then $\frac{(\ell - \frac{1}{2})^2}{\ell + 4} \geq \frac{\ell}{20}$. Therefore,

$$H(X) \leq n - \frac{(1-\epsilon)^2 \ell}{320}. \quad \square$$

6.2 Proof of the List-Decoding Direct Product Theorem

The proof of [Theorem 2](#) relies on one more technical observation.

Proposition 19. *For any constant $\lambda > 0$, if $k \in \mathbb{N}$ satisfies*

$$\frac{2^{-\frac{1}{2k}} - 2^{-\frac{1}{2}}}{1 - 2^{-\frac{1}{2}}} \geq \lambda$$

and $\epsilon \in (0, \frac{1}{2k}]$, then for every $\ell \in \mathbb{N}$,

$$\frac{2^{-\epsilon\ell} - 2^{-k\epsilon\ell}}{1 - 2^{-k\epsilon\ell}} \geq \lambda^\ell.$$

Proof. As a first step, define the function

$$g: y \mapsto (k-1) - ky + y^k.$$

This function satisfies $g(1) = 0$. Its derivative is $g'(y) = -k + ky^{k-1}$, which is negative on all $y \in [0, 1)$ and positive on all $y > 1$. So the function g is non-negative in the entire range $[0, \infty)$.

Consider now the function $h: \epsilon \mapsto \frac{2^{-\epsilon} - 2^{-k\epsilon}}{1 - 2^{-k\epsilon}}$. The derivative of this function is

$$h'(\epsilon) = \frac{2^{-(k+1)\epsilon} \ln 2}{(1 - 2^{-k\epsilon})^2} \cdot (-(k-1) + k2^\epsilon - (2^\epsilon)^k) = -C_{k,\epsilon} \cdot g(2^\epsilon)$$

with $C_{k,\epsilon} = \frac{2^{-(k+1)\epsilon} \ln 2}{(1 - 2^{-k\epsilon})^2} > 0$ for all $\epsilon > 0$. By the non-negativity of g , we then have that h is non-increasing on $[0, 1]$ and so every $\epsilon \in (0, \frac{1}{2k}]$ satisfies

$$\frac{2^{-\epsilon} - 2^{-k\epsilon}}{1 - 2^{-k\epsilon}} \geq \frac{2^{-\frac{1}{2k}} - 2^{-\frac{1}{2}}}{1 - 2^{-\frac{1}{2}}} \geq \lambda.$$

To complete the proof, we now want to show that for every $a \in (0, 1)$ and every $\ell \geq 1$,

$$\frac{a^\ell - a^{k\ell}}{1 - a^{k\ell}} \geq \left(\frac{a - a^k}{1 - a^k} \right)^\ell. \quad (9)$$

Establishing this inequality completes the proof of the proposition by taking $a = 2^{-\epsilon}$ and using the lower bound $\frac{2^{-\epsilon} - 2^{-k\epsilon}}{1 - 2^{-k\epsilon}} \geq \lambda$ established above.

Fix any $a \in (0, 1)$ and $k \in \mathbb{N}$. Consider the function

$$\phi: \ell \mapsto \ln \left(\frac{a^\ell - a^{k\ell}}{1 - a^{k\ell}} \right) - \ell \ln \left(\frac{a - a^k}{1 - a^k} \right).$$

To complete the proof of (9), we want to show that $\phi(\ell) \geq 0$ for all $\ell \in \mathbb{N}$. We already have $\phi(1) = 0$. So we want to show that $\phi'(\ell) \geq 0$ for all $\ell \geq 1$. The derivative of ϕ is

$$\phi'(\ell) = \ln(a) \left(\frac{a^\ell - ka^{k\ell}}{a^\ell - a^{k\ell}} + \frac{ka^\ell}{1 - a^{k\ell}} \right) - \ln \left(\frac{a - a^k}{1 - a^k} \right).$$

Using the bound $\frac{a - a^k}{1 - a^k} < a$ that holds for all $a \in (0, 1)$ and $k > 1$, we have that

$$\phi'(\ell) \geq \ln(a) \left(\frac{a^\ell - ka^{k\ell}}{a^\ell - a^{k\ell}} + \frac{ka^\ell}{1 - a^{k\ell}} - 1 \right).$$

Since $\ln(a) < 0$, we obtain our proof that $\phi'(\ell) \geq 0$ if we can show that

$$\frac{a^\ell - ka^{k\ell}}{a^\ell - a^{k\ell}} + \frac{ka^\ell}{1 - a^{k\ell}} \leq 1.$$

Expanding, re-arranging the terms, and simplifying, this inequality is equivalent to showing that

$$0 \leq a^k((k - 1) - ka + a^k) = a^k \cdot g(a).$$

As we have already seen, g is non-negative and so we also have that $\phi'(\ell) \geq 0$, and as a result (9) holds for all $\ell \in \mathbb{N}$, completing the proof of the proposition. \square

We are now ready to complete the proof of [Theorem 2](#).

Theorem 2 (List-Decoding Direct Product Theorem). *There exists a universal constant $c \in (\frac{1}{2}, 1)$ such that for every (possibly partial) function $f: \{0, 1\}^m \rightarrow \{0, 1\}$, $n \geq 1$, $\ell \in \{1, 2, \dots, n\}$, and $\gamma \in [c, 1)$,*

$$R_{\gamma^\ell}(\text{LD}_\ell^n \circ f) = \Omega(\ell R(f)).$$

Proof. Let us first introduce and fix a few technical parameters. Set $\lambda = 2^{-\frac{1}{2560}}$. Choose k to be the least positive integer that satisfies the inequality $\frac{2^{-1/2k} - 2^{-1/2}}{1 - 2^{-1/2}} \geq \lambda$. Let ϵ be defined so that $\gamma = 2^{-\epsilon}$.

Fix the universal constant c in the lemma statement to be $c = 2^{-\frac{1}{2k}}$ so that $\epsilon \leq \frac{1}{2k}$ when $\gamma \geq c$.

Fix any distribution μ on the domain of f . Let A be a deterministic algorithm that computes $\text{LD}_\ell^n \circ f$ with a success probability of at least $2^{-\epsilon\ell}$ on inputs drawn from μ^n . Let us define $\beta = 2^{-k\epsilon\ell}$ and $\alpha = \frac{2^{-\epsilon\ell} - \beta}{1 - \beta}$. The algorithm A has probability at least α of landing in a leaf with success probability at least β since our choice of α and β satisfies

$$\alpha \cdot 1 + (1 - \alpha)\beta = 2^{-\epsilon\ell}.$$

By [Lemma 18](#), the entropy loss score of each leaf with success probability at least $\beta = 2^{-k\epsilon\ell}$ is at least $2^{(1-k\epsilon)^2\ell/320-n}$. Therefore, the overall entropy loss score of A is at least $\alpha \cdot 2^{(1-k\epsilon)^2\ell/320-n}$. By the choice of k and [Proposition 19](#), we have $\alpha \geq \lambda^\ell = 2^{-\ell/2560} \geq 2^{-(1-k\epsilon)^2\ell/640}$. So the overall entropy score of A is at least $2^{(1-k\epsilon)^2\ell/640-n}$.

The argument above implies that

$$R_{2^{-\epsilon\ell}}(\text{LD}_\ell^n \circ f) \geq R_{\text{Ent } 2^{(1-k\epsilon)^2\ell/640-n}}^{\mu^n}(f^n).$$

By the Direct Product Theorem with the entropy loss score, we also have the bound

$$R_{\text{Ent } 2^{(1-k\epsilon)^2\ell/640-n}}^{\mu^n}(f^n) \geq \Omega\left(n \cdot \overline{\text{sR}}_{\text{Ent } 2^{(1-k\epsilon)^2\ell n/640-1}}^\mu(f)\right).$$

Since the success of any leaf of a randomized algorithm for f is at least $\frac{1}{2}$,

$$\overline{\text{sR}}_{\text{Ent } 2^{(1-k\epsilon)^2\ell n/640-1}}^\mu(f) \geq \overline{\text{R}}_{\text{Ent } 2^{(1-k\epsilon)^2\ell n/640-1}}^\mu(f).$$

Putting together the inequalities above and choosing μ that maximizes $\overline{\text{R}}_{\text{Ent } 2^{\alpha-1}}^\mu(f)$, we can then apply [Lemma 22](#) to obtain

$$R_{2^{-\epsilon\ell}}(\text{LD}_\ell^n \circ f) = \Omega\left((1-k\epsilon)^2\ell \cdot R_{\frac{2}{3}}(f)\right) = \Omega(\ell R(f)). \quad \square$$

6.3 Proof of the Threshold Direct Product Theorem

Proof of [Corollary 3](#). Let A be an algorithm that computes $\text{Thr}_k^n \circ f$ with success probability at least $\gamma^{\tau_{n,k}}$. Let $y \in \{0,1\}^n$ denote the output of A on some input $x \in \{0,1\}^{m \times n}$. Then the set

$$S_y = \{z \in \{0,1\}^n : \sum_{i=1}^n \mathbf{1}[z_i \neq y_i] \leq n-k\}$$

includes the true value $z^* = (f(x^{(1)}), \dots, f(x^{(n)}))$ if and only if A correctly computed $\text{Thr}_k^n \circ f$ on its execution.

Let B be the algorithm that simulates A and outputs the set S_y corresponding to A 's output y . The size of the list that B outputs is $|S_y| = 2^{n-\tau_{n,k}}$. And it computes $\text{LD}_{\tau_{n,k}}^n \circ f$ correctly with success probability $\gamma^{\tau_{n,k}}$. Therefore, by [Theorem 2](#), the algorithm A simulated by B must have query complexity at least $\Omega(\tau_{n,k} R(f))$. \square

6.4 Proof of the Labelled Threshold DPT

Proof of [Corollary 4](#). Let A compute $\text{Label}_k^n \circ f$ with success probability γ^k . Consider the algorithm B that simulates A and that on A 's output $y \in \{0,1,*\}^n$ returns the set

$$S_y = \{z \in \{0,1\}^n : \forall i \in [n], y_i \in \{z_i, *\}\}.$$

This set has cardinality $|S_y| = 2^{n-k}$ and it includes the true value $(f(x^{(1)}), \dots, f(x^{(n)}))$ if and only if $y \in \text{Label}_k^n \circ f$. Therefore, B succeeds in computing $\text{LD}_k^n \circ f$ with probability at least γ^k and so by [Theorem 2](#), the simulated algorithm A must have query complexity $\Omega(k R(f))$. \square

7 Score Amplification

7.1 Success Amplification by Boosting

Efficient success amplification is one of the key tools in the study of worst-case and average-case randomized complexity measures. It is not a property shared by distributional complexity when we fix a single distribution. (See [Sha04; BKST24] for example.) But, interestingly and critically for the proof of the upper bound in [Theorem 1](#), the maximum distributional complexity of functions *does* admit efficient success amplification.

Lemma 20. *For every success probability γ in the range $\frac{1}{2} \leq \gamma \leq 1$ and every Boolean-valued function $f: \mathcal{X} \rightarrow \{0, 1\}$,*

$$\overline{\mathbf{R}}_{3\gamma^2-2\gamma^3}(f) \leq 4 \cdot \overline{\mathbf{R}}_\gamma(f).$$

Proof. The proof is a direct extension of Schapire's original boosting algorithm [Sch90].

Fix any distribution μ on \mathcal{X} . We want to show that $\overline{\mathbf{R}}_{3\gamma^2-2\gamma^3}^\mu \leq 4 \cdot \overline{\mathbf{R}}_\gamma(f)$. We will select three randomized algorithms A_1 , A_2 , and A_3 . For each $i = 1, 2, 3$ and $x \in \mathcal{X}$, let $\alpha_i(x)$ denote the probability that A_i outputs the correct value $f(x)$.

Let A_1 be a randomized algorithm with average success probability $\sum_{x \in \mathcal{X}} \mu(x) \alpha_1(x) = \gamma$ and average cost $\overline{\text{cost}}^\mu(A_1) = \overline{\mathbf{R}}_\gamma^\mu(f)$ against μ .

Define ν to be the distribution on \mathcal{X} obtained by setting

$$\nu(x) = \mu(x) \left(\frac{(1 - \gamma)\alpha_1(x) + \gamma(1 - \alpha_1(x))}{2\gamma(1 - \gamma)} \right)$$

for every $x \in \mathcal{X}$. We can verify that ν is indeed a valid distribution since $\nu(x) \geq 0$ for all x and

$$\sum_x \nu(x) = \frac{\sum_x \mu(x) \alpha_1(x)}{2\gamma} + \frac{\sum_x \mu(x) (1 - \alpha_1(x))}{2(1 - \gamma)} = 1.$$

We select A_2 to be a randomized algorithm with average success probability $\sum_{x \in \mathcal{X}} \nu(x) \alpha_2(x) = \gamma$ and average cost $\overline{\text{cost}}^\nu(A_2) = \overline{\mathbf{R}}_\gamma^\nu(f)$ against ν .

Define π to be the distribution on \mathcal{X} obtained by setting

$$\pi(x) = \mu(x) \left(\frac{\alpha_1(x)(1 - \alpha_2(x)) + (1 - \alpha_1(x))\alpha_2(x)}{\Pr_{x \sim \mu}[A_1(x) \neq A_2(x)]} \right).$$

We select A_3 to be a randomized algorithm with average success probability $\sum_{x \in \mathcal{X}} \pi(x) \alpha_3(x) = \gamma$ and average cost $\overline{\text{cost}}^\pi(A_3) = \overline{\mathbf{R}}_\gamma^\pi(f)$ against π .

We combine the algorithms A_1 , A_2 , and A_3 to obtain an algorithm B for f that proceeds as follows. On input $x \in \mathcal{X}$, the algorithm B :

1. Runs A_1 on x , obtains the value y_1 .
2. Runs A_2 on x , obtains the value y_2 .
3. If $y_1 = y_2$, B outputs the same value.
4. Otherwise, B runs A_3 on x and outputs its value.

We now analyze the correctness and the cost of the algorithm B .

Correctness analysis. We want to show that B has average success probability at least $3\gamma^2 - 2\gamma^3$ when x is drawn from μ .

The probability that $A_1(x) \neq A_2(x)$ and that $A_3(x) = f(x)$ is

$$\sum_{x \in \mathcal{X}} \mu(x) (\alpha_1(x)(1 - \alpha_2(x)) + (1 - \alpha_1(x))\alpha_2(x)) \alpha_3(x) = \Pr_{x \sim \mu} [A_1(x) \neq A_2(x)] \cdot \sum_x \pi(x) \alpha_3(x).$$

By the correctness guarantee on A_3 , $\sum_x \pi(x) \alpha_3(x) = \gamma$. And

$$\Pr_{x \sim \mu} [A_1(x) \neq A_2(x)] = \sum_{x \in \mathcal{X}} \mu(x) (\alpha_1(x)(1 - \alpha_2(x)) + (1 - \alpha_1(x))\alpha_2(x)).$$

The guarantee on A_1 implies that $\sum_x \mu(x) \alpha_1(x) = \gamma$ so that

$$\Pr_{x \sim \mu} [A_1(x) \neq A_2(x) \wedge A_3(x) = f(x)] = \gamma^2 + \gamma \sum_x \mu(x) ((1 - \alpha_1(x))\alpha_2(x) - \alpha_1(x)\alpha_2(x)).$$

Also,

$$\Pr_{x \sim \mu} [A_1(x) = A_2(x) = f(x)] = \sum_{x \in \mathcal{X}} \mu(x) \alpha_1(x) \alpha_2(x).$$

So the overall probability that B outputs the correct value $f(x)$ is

$$\begin{aligned} \Pr_{x \sim \mu} [B(x) = f(x)] &= \Pr_{x \sim \mu} [A_1(x) = A_2(x) = f(x)] + \Pr_{x \sim \mu} [A_1(x) \neq A_2(x) \wedge A_3(x) = f(x)] \\ &= \gamma^2 + \sum_x \mu(x) ((1 - \gamma)\alpha_1(x) + \gamma(1 - \alpha_1(x))) \alpha_2(x) \\ &= \gamma^2 + 2\gamma(1 - \gamma) \sum_x \nu(x) \alpha_2(x) \end{aligned}$$

and the correctness guarantee on A_2 implies that $\sum_x \nu(x) \alpha_2(x) = \gamma$ which means that the overall success probability is $\gamma^2 + 2\gamma^2(1 - \gamma) = 3\gamma^2 - 2\gamma^3$, as we wanted to show.

Cost analysis. The average cost of A_1 against μ is bounded by

$$\overline{\text{cost}}^\mu(A_1) = \overline{\mathbf{R}}_\gamma^\mu(f) \leq \overline{\mathbf{R}}_\gamma(f).$$

Let $c_2(x)$ denote the average cost of A_2 on input x . The average cost of A_2 with respect to μ is

$$\overline{\text{cost}}^\mu(A_2) = \sum_x \mu(x) c_2(x) = \sum_x \frac{2}{\alpha_1(x)/\gamma + (1 - \alpha_1(x))/(1 - \gamma)} \nu(x) c_2(x).$$

For any input x , $0 \leq \alpha_1(x) \leq 1$ and so

$$\frac{2}{\alpha_1(x)/\gamma + (1 - \alpha_1(x))/(1 - \gamma)} \leq 2 \max\{\gamma, 1 - \gamma\} \leq 2.$$

Also, the cost guarantee on A_2 implies that $\sum_x \nu(x) c_2(x) \leq \overline{\mathbf{R}}_\gamma(f)$ so

$$\overline{\text{cost}}^\mu(A_2) \leq 2 \sum_x \nu(x) c_2(x) \leq 2\overline{\mathbf{R}}_\gamma(f).$$

Finally, for every $x \in \mathcal{X}$, the algorithm B runs A_3 with probability $\alpha_1(x)(1 - \alpha_2(x)) + (1 - \alpha_1(x))\alpha_2(x)$. So if we let $c_3(x)$ denote the expected cost of A_3 on input x , the average contribution of A_3 to the overall cost of B is

$$\sum_x \mu(x)(\alpha_1(x)(1 - \alpha_2(x)) + (1 - \alpha_1(x))\alpha_2(x))c_3(x) = \Pr_{x \sim \mu}[A_1(x) \neq A_2(x)] \sum_x \pi(x)c_3(x).$$

The cost guarantee on A_3 implies that $\sum_x \pi(x)c_3(x) \leq \bar{R}_\gamma(f)$. So its overall contribution to the cost of B is at most $\Pr_{x \sim \mu}[A_1(x) \neq A_2(x)] \cdot \bar{R}_\gamma(f) \leq \bar{R}_\gamma(f)$.

In total, the average cost of B over μ is thus bounded above by

$$\overline{\text{cost}}^\mu(B) \leq \bar{R}_\gamma(f) + 2\bar{R}_\gamma(f) + \bar{R}_\gamma(f) = 4\bar{R}_\gamma(f). \quad \square$$

The boosting lemma immediately yields the score amplification result stated in the preliminaries, which we restate here for convenience.

Lemma 8. *For any $\delta > 0$, any $f: \{0, 1\}^m \rightarrow \{0, 1\}$, and any success probability parameter $\gamma \in [\frac{1+\delta}{2}, 1 - \delta]$,*

$$\bar{R}_{\gamma + \frac{\delta}{3}}(f) \leq 4 \cdot \bar{R}_\gamma(f).$$

Proof. This lemma follows immediately from the last one by observing that the difference in the success probability $3\gamma^2 - 2\gamma^3$ after boosting and the original success probability γ is

$$3\gamma^2 - 2\gamma^3 - 1 = (2\gamma - 1)(1 - \gamma).$$

When $\frac{1+\delta}{2} \leq \gamma \leq \frac{2}{3}$, this difference is bounded below by $(2\gamma - 1)(1 - \gamma) \geq (2(\frac{1+\delta}{2}) - 1)(1 - \frac{2}{3}) = \frac{\delta}{3}$. And when $\frac{2}{3} \leq \gamma \leq 1 - \delta$, we have $(2\gamma - 1)(1 - \gamma) \geq (2(\frac{2}{3}) - 1)(1 - (1 - \delta)) = \frac{\delta}{3}$. \square

7.2 Hellinger Score Amplification

The success amplification result from the last section also shows that for every small δ , $R_{\frac{2}{3}}(f) = \Theta(\bar{R}_{\frac{2}{3}}(f)) = \Theta(\frac{1}{\delta^2} \bar{R}_{\frac{1+\delta}{2}}(f))$. When we use Hellinger score, we can get more efficient success amplification where the multiplicative factor is linear instead of quadratic in $\frac{1}{\delta}$.

The technical lemma at the heart of the proof is stated more naturally in terms of *Hellinger loss* instead of Hellinger score. Define the Hellinger loss of a leaf ℓ with respect to f and μ to be $\text{Hel}_{\downarrow f}^\mu(\ell) = 2\text{Hel}_f^\mu(\ell) - 1$. The Hellinger loss of a randomized algorithm is the expected Hellinger loss of its leaves: $\overline{\text{Hel}}_{\downarrow f}^\mu(R) = \mathbb{E}_{D \sim R} \mathbb{E}_{\ell \sim D(\mu)} \text{Hel}_{\downarrow f}^\mu(\ell)$. We write $\bar{R}_{\text{Hel}_{\downarrow \epsilon}}(f)$ to denote the maximum distributional query complexity of f for the set of algorithms with Hellinger loss at most ϵ .

Lemma 21. *For any $\alpha \in [\frac{1}{2}, 1]$, $\beta \in [0, 1]$ and function $f: \{0, 1\}^m \rightarrow \{0, 1\}$,*

$$\bar{R}_{\text{Hel}_{\downarrow \alpha(1 - \frac{1-\beta}{6})}}(f) \leq \bar{R}_{\text{Hel}_{\downarrow \alpha}}(f) + 2\bar{R}_{\text{Hel}_{\downarrow \beta}}(f).$$

Proof. Fix any distribution μ on the inputs to f . Take A to be a randomized decision tree with Hellinger loss at most α and average cost at most $\bar{R}_{\text{Hel}_{\downarrow \alpha}}(f)$ with respect to μ .

For any set leaf ℓ of a decision tree in the support of A , define the distribution

$$\nu_\ell(x) = \begin{cases} \frac{\mu(x)}{2\mu(\ell_0)} & \text{if } x \in \ell_0 \\ \frac{\mu(x)}{2\mu(\ell_1)} & \text{if } x \in \ell_1 \\ 0 & \text{if } x \notin \ell. \end{cases}$$

Let B_ℓ denote a randomized algorithm with Hellinger loss at most β and cost at most $\overline{\overline{R}}_{\text{Hel}_\downarrow \beta}(f)$ against the distribution ν_ℓ .

Consider now the randomized algorithm C defined by the following procedure: First, we run A on the input x . Let ℓ denote the leaf of the tree reached by input x and the subset of inputs $\{0, 1\}^m$ that lead to this leaf. Then we run B_ℓ on x .

The Hellinger loss of the algorithm C is bounded by

$$\begin{aligned} \overline{\text{Hel}}_\downarrow^\mu(C) &= \mathbb{E}_{D \sim A} \sum_{\ell \in D} \mathbb{E}_{D' \sim B_\ell} \sum_{\ell' \in D'} 2\sqrt{\mu(\ell'_0)\mu(\ell'_1)} \\ &= \mathbb{E}_{D \sim A} \sum_{\ell \in D} \mathbb{E}_{D' \sim B_\ell} \sum_{\ell' \in D'} 2\sqrt{\mu(\ell_0)\mu(\ell_1)} \cdot 2\sqrt{\nu_\ell(\ell'_0)\nu_\ell(\ell'_1)} \\ &= \mathbb{E}_{D \sim A} \sum_{\ell \in D} 2\sqrt{\mu(\ell_0)\mu(\ell_1)} \cdot \left(\mathbb{E}_{D' \sim B_\ell} \sum_{\ell' \in D'} 2\sqrt{\nu_\ell(\ell'_0)\nu_\ell(\ell'_1)} \right) \\ &\leq \alpha \cdot \beta. \end{aligned}$$

By construction, the expected cost of A with respect to μ is at most $\overline{\overline{R}}_{\text{Hel}_\downarrow \alpha}(f)$. And for any set ℓ , the expected cost of B_ℓ on inputs drawn from μ conditioned on being in ℓ is

$$\begin{aligned} \sum_{x \in \ell} \frac{\mu(x)}{\mu(\ell)} \overline{\text{cost}}(B_\ell, x) &= \sum_{x \in \ell_0} \frac{2\mu(\ell_0)}{\mu(\ell)} \nu_\ell(x) \overline{\text{cost}}(B_\ell, x) + \sum_{x \in \ell_1} \frac{2\mu(\ell_1)}{\mu(\ell)} \nu_\ell(x) \overline{\text{cost}}(B_\ell, x) \\ &\leq 2 \frac{\max\{\mu(\ell_0), \mu(\ell_1)\}}{\mu(\ell)} \sum_x \nu_\ell(x) \overline{\text{cost}}(B_\ell, x). \end{aligned}$$

By the trivial upper bound $\max\{\mu(\ell_0), \mu(\ell_1)\} \leq \mu(\ell)$ and the guarantee on the expected cost of B_ℓ under the distribution ν_ℓ , we then have that $\overline{\text{cost}}(B_\ell, \mu|_\ell) \leq 2 \overline{\overline{R}}_\beta(f)$. Therefore, the overall expected cost of B is at most $\overline{\overline{R}}_{\text{Hel}_\downarrow \alpha}(f) + 2 \overline{\overline{R}}_{\text{Hel}_\downarrow \beta}(f)$, as claimed. \square

We then get the linear amplification lemma stated in the preliminaries as a corollary.

Lemma 9. *For any $\delta \in [0, 1]$ and any $f: \{0, 1\}^m \rightarrow \{0, 1\}$,*

$$R(f) = O\left(\frac{1}{\delta} \cdot \overline{\overline{R}}_{\text{Hel } \frac{1+\delta}{2}}(f)\right).$$

Proof. By repeated application of [Lemma 21](#) with the initial parameter $\alpha = 1$ and the observation that $\overline{\overline{R}}_{\text{Hel}_\downarrow 1}(f) = 0$ for all functions f , we have that for any $k \geq 1$ and $\beta \in [0, 1]$,

$$\overline{\overline{R}}_{\text{Hel}_\downarrow \beta^k}(f) \leq 2k \cdot \overline{\overline{R}}_{\text{Hel}_\downarrow \beta}(f).$$

Therefore, taking $k = \lceil \frac{2}{1-\beta} \rceil$, we have that

$$\overline{\overline{R}}_{\text{Hel}_\downarrow e^{-1}}(f) = O\left(\frac{1}{1-\beta} \overline{\overline{R}}_{\text{Hel}_\downarrow \beta}(f)\right).$$

Since Hellinger score $\frac{1+\delta}{2}$ is equivalent to Hellinger loss $1 - \delta$, the above identity implies that

$$\overline{\overline{R}}_{\text{Hel}_\downarrow e^{-1}}(f) = O\left(\frac{1}{\delta} \cdot \overline{\overline{R}}_{\text{Hel } \frac{1+\delta}{2}}(f)\right).$$

The lemma then follows from observing that $R(f) = O(\overline{\overline{R}}_{\text{Hel}_\downarrow c}(f))$ for any constant c bounded away from 0 and 1. \square

The linear amplification lemma for Hellinger score also immediately implies an analogous lemma for exponential entropy score.

Lemma 22. *For any α , $0 \leq \alpha \leq 1$,*

$$R(f) = O\left(\frac{1}{\alpha} \cdot \overline{R}_{\text{Ent } 2^{\alpha-1}}(f)\right).$$

Proof. For any $x \in [0, 1]$, the convexity of the function 2^{-x} implies that $2^{-x} \leq 1 - \frac{x}{2}$. As a result, the exponential entropy score of a leaf with density q is

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

The function $1 - h(q)$ is in turn bounded above by

$$1 - h(q) \leq \frac{1}{\ln 2}(1 - 2\sqrt{q(1-q)}).$$

Consider any randomized algorithm R with exponential entropy score at least $2^{\alpha-1}$ with respect to f and a fixed distribution μ . Then with the two inequalities above,

$$2^{\alpha-1} \leq \overline{\text{Ent}}_f^\mu(R) = \mathbb{E}_\ell 2^{-h(\mu(\ell_1)/\mu(\ell))} \leq \frac{1 + (\ln 2)^{-1} \mathbb{E}_\ell [1 - 2\sqrt{\frac{\mu(\ell_1)}{\mu(\ell)} \cdot \frac{\mu(\ell_0)}{\mu(\ell)}}]}{2}.$$

This inequality and the basic inequality $1 + \ln(2)\alpha \leq 2^\alpha$ imply

$$\ln(2)^2 \alpha \leq \mathbb{E}_\ell \left[1 - 2\sqrt{\frac{\mu(\ell_1)}{\mu(\ell)} \cdot \frac{\mu(\ell_0)}{\mu(\ell)}} \right].$$

And the Hellinger score of R is therefore bounded below by

$$\overline{\text{Hel}}_f^\mu(R) = \mathbb{E}_\ell \left[1 - \sqrt{\frac{\mu(\ell_1)}{\mu(\ell)} \cdot \frac{\mu(\ell_0)}{\mu(\ell)}} \right] = \frac{1 + \mathbb{E}_\ell [1 - 2\sqrt{\frac{\mu(\ell_1)}{\mu(\ell)} \cdot \frac{\mu(\ell_0)}{\mu(\ell)}}]}{2} \geq \frac{1 + \ln(2)^2 \alpha}{2}.$$

We have just shown that for every distribution μ , every randomized algorithm R with entropy score $\overline{\text{Ent}}_f^\mu(R) \geq 2^{\alpha-1}$ also has Hellinger score $\overline{\text{Hel}}_f^\mu(R) \geq \frac{1 + \ln(2)^2 \alpha}{2}$. As a result, $\overline{R}_{\text{Ent } 2^{\alpha-1}}(f) \leq \overline{R}_{\text{Hel } \frac{1 + \ln(2)^2 \alpha}{2}}(f)$ and the lemma is an immediate consequence of [Lemma 9](#). \square

8 Relations Between the Different Complexity Measures

We prove various relations between different complexity measures in this section.

8.1 Worst-Case and Average-Case Complexity

Relations between worst-case and average-case complexity measures are well known. We include them here for the convenience of the reader.

As a starting point, by definition, the worst-case and average-case randomized query complexity measures always satisfy the inequality

$$\overline{R}_\gamma(f) \leq R_\gamma(f).$$

In the other direction, we have the following inequalities.

Proposition 23. *For every $\frac{1}{2} < \gamma < 1$ and every function f ,*

$$R_\gamma(f) \leq O\left(\frac{1}{\min\{\gamma - \frac{1}{2}, 1 - \gamma\}} \cdot \bar{R}_\gamma(f)\right).$$

Proof. Consider first the small-advantage regime where $\gamma = \frac{1+\delta}{2}$. Let A be a randomized algorithm with success probability γ and average cost $\bar{R}_\gamma(f)$. Let A' be the algorithm obtained by truncating A and forcing it to guess the value of f whenever A exceeds its expected cost by a multiplicative factor of $4/\delta$. By Markov's inequality, the probability that this truncation event occurs is at most $\delta/4$, so that the success probability of A' is at least $\frac{1+\delta/2}{2}$ and we obtain the desired bound on the worst-case complexity for success probability $\frac{1+\delta}{2}$ by applying standard success amplification.

For the small-error regime where $\gamma = 1 - \epsilon$, we use the same argument with truncation factor $1/\epsilon$. \square

Proposition 23 implies that the average-case and worst-case randomized query complexity measures are asymptotically equivalent in the bounded-error regime, when γ is bounded away from both $\frac{1}{2}$ and 1 by a constant. In the small-bias and small-error regimes, however, there are functions whose average-case randomized query complexity is asymptotically smaller than their worst-case randomized query complexity. For the small-bias regime, the bound in the last proposition is optimal.

Proposition 24. *For every $n \in \mathbb{N}$, there exists a function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ such that for every $\delta = \delta(n) > 0$*

$$\bar{R}_{\frac{1+\delta}{2}}(f) \leq \delta \cdot R_{\frac{1+\delta}{2}}(f).$$

Proof. Consider the parity function $\oplus_n: x \mapsto \sum_{i=1}^n x_i \pmod{2}$. For any $\delta > 0$, $R_{\frac{1+\delta}{2}}(\oplus_n) = n$. But in the average-case complexity setting, the (lazy) algorithm that with probability δ queries everything to compute the value of $\oplus_n(x)$ and otherwise guesses the value of the function has average cost $\delta n = \delta R(\oplus_n)$ and succeeds with probability $\delta + (1 - \delta)\frac{1}{2} = \frac{1+\delta}{2}$ on every input. \square

Note that by contrast there are also other functions for which $\bar{R}_{\frac{1+\delta}{2}}(f) = \Theta(R_{\frac{1+\delta}{2}}(f))$. One example of such a function is the Majority function when $\delta = 1/n$.

There can also be asymptotic separations between the average-case and worst-case randomized query complexities in the small-error regime.

Proposition 25. *For infinitely many values of $n \in \mathbb{N}$, there exists a partial function $f: X \rightarrow \{0, 1\}$, $X \subseteq \{0, 1\}^n$, such that for every $\epsilon = \epsilon(n) > 0$,*

$$\bar{R}_{1-\epsilon}(f) \leq O\left(\frac{1}{\log(1/\epsilon)} \cdot R_{1-\epsilon}(f)\right).$$

Proof. Fix any even value of n and define the set

$$X = \left\{ x \in \{0, 1\}^n : \sum_{i=1}^{n/2} x_{2i-1} = \frac{n}{2} \wedge \exists b \in \{0, 1\} \forall i \in [\frac{n}{2}] x_{2i-1} = 1 \Rightarrow x_{2i} = b \right\}.$$

For each $x \in X$, we define $f(x)$ to be the value of b for which the last condition is satisfied.

Informally, we can think of f as containing $\frac{n}{2}$ bins. Half of the bins are useful in that the second bit of the bin determines the value of the function and the other half are useless. Furthermore, the first bit of each bit indicates whether the bin is useful or not.

An algorithm that computes f with error at most ϵ must find a useful bin with probability at least $1 - O(\epsilon)$. This can be guaranteed only when the worst-case query complexity of the algorithm is $\Omega(\log \frac{1}{\epsilon})$. So $R_{1-\epsilon}(f) = \Omega(\log \frac{1}{\epsilon})$.

But the average-case cost of the algorithm that examines bins uniformly at random follows a geometric distribution with success parameter $\frac{1}{2}$ and so $\overline{R}_{1-\epsilon}(f) = O(1)$. \square

There are also total functions that exhibit an analogous gap between the worst-case and average-case query complexity; see [BB19].

8.2 Maximum Distributional Complexity

Any algorithm with score γ and average cost at most c on every input x also has the same score and cost guarantees over all distributions μ on the inputs. So for every f and γ ,

$$\overline{R}_\gamma(f) \leq \overline{R}_\gamma(f).$$

This inequality is asymptotically optimal in the bounded-error regime. It is also asymptotically optimal in the small-error regime, where Vereshchagin [Ver98] showed that $\overline{R}_{1-\epsilon}(f) \leq 2\overline{R}_{1-\epsilon/2}(f)$.

In the small-advantage regime, we have the following relation between the maximum distributional and average-case complexity.

Proposition 26. *For every function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ and $\delta = \delta(n) > 0$, the average-case and maximum distributional complexity measures for success probability $\frac{1+\delta}{2}$ are related by*

$$\overline{R}_{\frac{1+\delta}{2}}(f) \leq O\left(\frac{1}{\delta} \cdot \overline{R}_{\frac{1+\delta}{2}}(f)\right).$$

Proof. We adapt Vereshchagin's argument [Ver98] for the result mentioned above to the small-advantage regime.

Fix any pair of distributions μ and ν . For every $0 \leq \lambda \leq 1$, let $\pi_\lambda = \lambda\mu + (1 - \lambda)\nu$. And let A_λ be a randomized algorithm with success at least $\frac{1+\delta}{2}$ and average cost at most $\overline{R}_{\frac{1+\delta}{2}}(f)$ when both are measured against μ . The success probability of A_λ satisfies

$$\overline{\text{success}}_f^{\pi_\lambda}(A_\lambda) = \lambda \cdot \overline{\text{success}}_f^\mu(A_\lambda) + (1 - \lambda)\overline{\text{success}}_f^\nu(A_\lambda).$$

Since $\overline{\text{success}}_f^\mu(A_\lambda) \leq 1$, the success of A_λ against ν is bounded below by

$$\overline{\text{success}}_f^\nu(A_\lambda) \geq \frac{(1 + \delta)/2 - \lambda}{1 - \lambda}.$$

Similarly, the average cost of A_λ satisfies

$$\overline{\text{cost}}^{\pi_\lambda}(A_\lambda) = \lambda \overline{\text{cost}}^\mu(A_\lambda) + (1 - \lambda)\overline{\text{cost}}^\nu(A_\lambda)$$

so that

$$\overline{\text{cost}}^\mu(A_\lambda) \leq \frac{1}{\lambda} \cdot \overline{R}_{\frac{1+\delta}{2}}(f).$$

By choosing $\lambda = \delta/4$, we obtain an algorithm A with success probability $\overline{\text{success}}_f^\nu(A) \geq \frac{1+\delta/2}{2}$ and average cost $\overline{\text{cost}}^\mu(A) \leq \frac{4}{\delta} \overline{R}_{\frac{1+\delta}{2}}(f)$. This means that

$$\overline{R}_{\frac{1+\delta/2}{2}}(f) \leq \frac{4}{\delta} \cdot \overline{R}_{\frac{1+\delta}{2}}(f)$$

and the result now follows from success amplification for average-case complexity. \square

The bound in [Proposition 26](#) is qualitatively tight in that there can be a polynomial factor of δ that separates the average-case and maximum distributional complexity of functions.

Proposition 27. *For infinitely many $n \in \mathbb{N}$, there exists a partial function $f: X \rightarrow \{0, 1\}$, $X \subseteq \{0, 1\}^n$ such that for all $\delta = \delta(n) > 0$,*

$$\overline{\overline{R}}_{\frac{1+\delta}{2}}(f) = O\left(\delta^{\frac{1}{6}} \overline{R}_{\frac{1+\delta}{2}}(f)\right).$$

Proof. Consider the following construction in which we split the n coordinates into k blocks B_1, \dots, B_k of n/k variables each. We construct the function so that there are two types of blocks: hard blocks and easy blocks. The first variable in each block determines the type of a block, so that the algorithm can tell if a block is hard or easy with a single query. We also associate a value with each block. For each easy block, the value is determined by the second variable in the block, so that the value of an easy block is determined with only one additional query. The value of a hard block, by contrast, is the parity of $1/\delta$ variables, so that this value can only be determined by making $1/\delta$ queries to the block.

Formally, writing $x \in \{0, 1\}^n$ as $x = (x^{(1)}, \dots, x^{(k)})$ with each $x^{(i)} \in \{0, 1\}^{n/k}$, we define the function $g: \{0, 1\}^{n/k} \rightarrow \{0, 1\}$ by

$$g(y) = \begin{cases} y_2 & \text{if } y_1 = 0 \\ y_2 \oplus \dots \oplus y_{1/\delta+1} & \text{if } y_1 = 1. \end{cases}$$

Then the function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ is

$$f(x) = \text{Maj}(g(x^{(1)}), \dots, g(x^{(k)})).$$

Let us consider the partial function obtained by restricting f 's domain X to one of two possible types of input:

Many hard blocks: For inputs of this type, there are exactly $\sqrt{\delta}k$ blocks that are hard. And their values are balanced: exactly half of them have value 0 and the other half have value 1. The values of the easy blocks have δ bias towards the value of the function. (I.e., for inputs of this type, the function f is the Gap Majority function on the values of the easy blocks.)

Few hard blocks: For inputs of this type, exactly δk blocks are hard. Each of these blocks has the same value, whereas the values of the easy blocks are balanced. As a result, the value of the function is the value of each of the hard blocks.

We now prove the two complexity bounds separately. First, we establish the upper bound

$$\overline{\overline{R}}_{\frac{1+\delta}{2}}(f) = O(1/\delta^3)$$

on the maximum distributional complexity of f . We do so by considering two natural algorithms for computing f . Algorithm A picks a random block. If the block is easy, it outputs its value; otherwise it just guesses. Algorithm B picks a random block. If the block is hard, it computes its value and outputs it; otherwise it just guesses.

Fix now any distribution μ on the domain of the partial function. Let p denote the probability that an input drawn from μ has many hard blocks.

Consider first the case where $p \geq \delta^{\frac{1}{6}}$. Algorithm A has success probability at least $\frac{1+p\delta}{2}$ on inputs drawn from μ . By extending the algorithm so that it queries the value of $O(1/p^2)$ easy blocks

instead of a single one (ignoring hard blocks, as before), we obtain an algorithm A' with success probability $\frac{1+\delta}{2}$ and average cost $O(1/p^2) \leq O(\delta^{-1/3})$ with respect to μ .

Consider now the case where $p < \delta^{\frac{1}{6}}$. Algorithm B has success probability at least $\frac{1+(1-p)\delta}{2}$ on inputs drawn from μ . By considering the extension of this algorithm that queries two blocks to see if either of them is hard, and outputting the value of a hard block if one is found, we obtain algorithm B' with success probability at least $\frac{1+(1-p)(2\delta-\delta^2)}{2}$ against μ , which is at least $\frac{1+\delta}{2}$ when $p \leq \delta^{\frac{1}{6}}$. And the average cost of B' over μ is

$$2 + p(2\sqrt{\delta} - \delta)(1/\delta) + (1-p)(2\delta - \delta^2)(1/\delta) \leq 2 + 2p/\sqrt{\delta} + 2 = O(p/\sqrt{\delta}),$$

which is $O(\delta^{-1/3})$ when $p < \delta^{\frac{1}{6}}$.

Combining the two observations above, we obtain that for any distribution μ on domain X of f , there is an algorithm (either A' or B') with success probability $\frac{1+\delta}{2}$ against μ and average cost $O(\delta^{-1/3})$ with respect to the same distribution. Therefore, by definition $O(\delta^{-1/3})$ is an upper bound on $\overline{R}_{\frac{1+\delta}{2}}(f)$.

To complete the proof of the proposition, we now show that the average-case complexity of f is bounded below by

$$\overline{R}_{\frac{1+\delta}{2}}(f) = \Omega(\delta^{-1/2}).$$

To do so, assume on the contrary that there is an algorithm that computes f with success probability at least $\frac{1+\delta}{2}$ and average cost $q = o(\sqrt{1/\delta})$. Any algorithm with this cost can only check $o(\sqrt{1/\delta})$ blocks to see if they are hard or easy and it can compute the value of at most that number of easy blocks. The goal of an algorithm seeking to minimize its average cost is to avoid computing the value of any hard block when there are many of them.

Without loss of generality, we can consider algorithms that decide whether to compute a hard block based on (i) the number of hard blocks observed; and (ii) the bias of the easy blocks that is computed. Note that when there are many hard blocks, the easy blocks have bias δ towards the function value, whereas when there are few hard blocks, the easy blocks's values are balanced. However, with at most $o(1/\sqrt{\delta})$ easy block values, an algorithm cannot distinguish between these two cases with bias more than $o(1)$. Therefore, without loss of generality we can consider algorithms that use only the number of hard blocks observed to determine whether they should compute the value of one of those blocks.

When we run an algorithm with query complexity $o(\sqrt{1/\delta})$ on inputs that have few hard blocks, the probability that the algorithm finds at least 1 block is at most $q \cdot \delta = o(\sqrt{\delta})$. And the probability that it finds more than one hard block is at most $\binom{q}{2}\delta^2 = o(\delta)$. So to have success probability at least $\frac{1+\delta}{2}$, its probability of computing a hard block (over all possible biases of the easy blocks) when there is exactly one such block must be at least $\sqrt{\delta}$. But then the probability that the algorithm sees exactly 1 hard block on inputs that have many hard blocks is constant. Which means that its average cost on the uniform distribution over those inputs is at least $\sqrt{\delta} \cdot 1/\delta = 1/\delta^2$. \square

We remark that the last proposition also implies a polynomial separation between average-case and maximum distributional complexity: there exists a partial function f and success parameter γ for which $\overline{R}_{\gamma}(f) = O(\overline{R}_{\gamma}(f)^{2/3})$. We have not tried to determine if the same is true for total function (though we suspect that it is) and what is the largest polynomial separation that we could achieve between the two complexity measures.

8.3 Score-Weighted Complexity

8.3.1 Relation to Maximum Distributional Complexity

We can give an upper bound as well as a lower bound of the score-weighted query complexity of randomized algorithms in terms of their maximum distributional complexity.

Proposition 10. *For every function $f: \{0, 1\}^m \rightarrow \{0, 1\}$ and score parameter γ ,*

$$\frac{1}{2} \cdot \overline{\mathbf{R}}_\gamma(f) \leq \overline{\mathbf{sR}}_\gamma(f) \leq \frac{1}{\gamma} \cdot \overline{\mathbf{R}}_\gamma(f)$$

Proof. For any distribution μ , we can express the score-weighted cost of f as

$$\overline{\mathbf{sR}}_\gamma^\mu(f) = \min_{R: \text{score}_f^\mu(R) \geq \gamma} \frac{\mathbb{E}_{\ell \sim L(\mu, R)}[\text{score}_f^\mu(\ell) \cdot \text{cost}(\ell)]}{\text{score}_f^\mu(R)}$$

where $L(\mu, R)$ the distribution on the leaves $\ell = D_l(x)$ obtained by drawing $D \sim R$ and $x \sim \mu$.

Since scoring measures are bounded above by 1,

$$\overline{\mathbf{sR}}_\gamma(f) = \max_\mu \overline{\mathbf{sR}}_\gamma^\mu(f) \leq \max_\mu \min_{R: \text{score}_f^\mu(R) \geq \gamma} \frac{\mathbb{E}_\ell[1 \cdot \text{cost}(\ell)]}{\gamma} = \frac{1}{\gamma} \cdot \overline{\mathbf{R}}_\gamma(f).$$

And since they are also bounded below by $\frac{1}{2}$,

$$\overline{\mathbf{R}}_\gamma(f) = \max_\mu \overline{\mathbf{R}}_\gamma^\mu(f) = 2 \cdot \max_\mu \min_{R: \text{score}_f^\mu(R) \geq \gamma} \frac{\mathbb{E}_\ell[\frac{1}{2} \cdot \text{cost}(\ell)]}{1} \leq 2 \cdot \overline{\mathbf{sR}}_\gamma(f). \quad \square$$

8.3.2 Relation to Worst-Case Query Complexity

We can also upper bound score-weighted complexity with worst-case complexity for every function and for their direct products.

Proposition 11. *For every function $f: \{0, 1\}^m \rightarrow \{0, 1\}$, parameter $n \in \mathbb{N}$, and score γ ,*

$$\overline{\mathbf{sR}}_{\gamma^n}(f) \leq \mathbf{R}_{\gamma^n}(f).$$

Proof. Fix any distribution μ on the domain of f^n . Let \mathcal{R}_μ denote the set of randomized algorithms with score at least γ^n on every input of f^n , maximum cost $\mathbf{R}_{\gamma^n}(f^n)$. The set \mathcal{R}_μ is non-empty since by definition there is an algorithm that satisfies these two properties. And every leaf ℓ of the decision trees in the support of the randomized algorithms in this set have score $\text{score}_{f^n}^\mu(\ell) \geq (\frac{1}{2})^n > 0$ by the boundedness of scoring measures.

For any $R \in \mathcal{R}_\mu$ and leaf ℓ of R ,

$$\text{score}_{f^n}^\mu(\ell) \cdot \text{cost}(\ell) \leq \text{score}_{f^n}^\mu(\ell) \cdot \mathbf{R}_{\gamma^n}(f^n).$$

So

$$\overline{\mathbf{sR}}_\gamma^\mu(R) = \frac{\mathbb{E}_\ell[\text{score}_{f^n}^\mu(\ell) \cdot \text{cost}(\ell)]}{\text{score}_{f^n}^\mu(R)} \leq \frac{\mathbb{E}_\ell[\text{score}_{f^n}^\mu(\ell) \cdot \mathbf{R}_{\gamma^n}(f^n)]}{\text{score}_{f^n}^\mu(R)} = \mathbf{R}_{\gamma^n}(f^n)$$

and $\overline{\mathbf{sR}}_{\gamma^n}(f^n) = \max_\mu \overline{\mathbf{sR}}_{\gamma^n}^\mu(f^n) \leq \max_\mu \min_{R \in \mathcal{R}_\mu} \overline{\mathbf{sR}}_\gamma^\mu(R) \leq \mathbf{R}_{\gamma^n}(f^n). \quad \square$

8.4 Score and Discounted Score

One of the basic properties of the discounted score measure that we will require in the proof of [Lemma 6](#) is that for any score parameter γ in the non-trivial range between the score of the guessing algorithm and 1, there exists some discount rate α^* for which the maximum discounted score $\text{DS}_{\alpha^*}^\mu(f)$ is achieved by a randomized algorithm R with score exactly equal to γ .

Proposition 28. *Fix any function f and distribution μ . For every success parameter γ that is at least as large as the success probability of the trivial guessing algorithm, there exists a discount rate α^* and a randomized algorithm R^* that satisfies*

$$\text{ds}_{f,\alpha^*}^\mu(R^*) = \text{DS}_{\alpha^*}^\mu(f) \quad \text{and} \quad \overline{\text{score}}_f^\mu(R^*) = \gamma.$$

Proof. Note first that since the discounted score of a randomized decision tree R is by definition

$$\text{ds}_{f,\alpha}^\mu(R) = \mathbb{E}_{D \sim R} \mathbb{E}_{x \sim \mu} [\text{score}_f^\mu(D(x)) \cdot e^{-\alpha \text{cost}(D(x))}],$$

there is always a deterministic decision tree D in the support of R with discounted score at least as large as that of R . As a result, the maximum discounted score $\text{DS}_\alpha^\mu(f)$ is always attained by a deterministic decision tree.

For each deterministic decision tree T , define the function $\phi_T: \alpha \mapsto \text{ds}_{f,\alpha}^\mu(T)$. The function ϕ_T is continuous. Then for any two trees T_1 and T_2 , the function $\phi_{T_1} - \phi_{T_2}$ is also continuous. As a result, the set $\{\alpha : \text{ds}_{f,\alpha}^\mu(T_1) \geq \text{ds}_{f,\alpha}^\mu(T_2)\}$, which is the preimage of $\phi_{T_1} - \phi_{T_2}$ on the closed set $[0, \infty)$, is a closed set.

Let us write $S_{\geq}, S_{\leq} \subseteq [0, \infty)$ to represent the sets of discount rates for which the maximum discount score is achieved by a deterministic algorithm with score at least and at most γ , respectively. The set S_{\geq} is closed, since we can write it as the finite union of finite intersections of closed sets

$$S_{\geq} = \bigcup_{T_1: \overline{\text{score}}_f^\mu(T_1) \geq \gamma} \bigcap_{T_2} \{\alpha : \text{ds}_{f,\alpha}^\mu(T_1) \geq \text{ds}_{f,\alpha}^\mu(T_2)\}.$$

Similarly, S_{\leq} is also closed. And these two closed sets partition their universe, i.e., $S_{\geq} \cup S_{\leq} = [0, \infty)$. As a result, their intersection $S_{\geq} \cap S_{\leq}$ is not empty. Let α^* be any element in this intersection.

By construction, for our choice of α^* there exist two deterministic algorithms T_1 and T_2 that satisfy

$$\text{ds}_{f,\alpha^*}^\mu(T_1) = \text{ds}_{f,\alpha^*}^\mu(T_2) = \text{DS}_{\alpha^*}^\mu(f)$$

and

$$\overline{\text{score}}_f^\mu(T_1) = \gamma - \delta_1, \quad \overline{\text{score}}_f^\mu(T_2) = \gamma + \delta_2$$

for some parameters $\delta_1, \delta_2 \geq 0$. Define the randomized algorithm A to be the convex combination

$$A = \frac{\delta_2}{\delta_1 + \delta_2} T_1 + \frac{\delta_1}{\delta_1 + \delta_2} T_2.$$

The discounted score of A is

$$\text{ds}_{f,\alpha^*}^\mu(A) = \text{DS}_{\alpha^*}^\mu(f)$$

and its score is

$$\overline{\text{score}}_f^\mu(A) = \frac{\delta_2}{\delta_1 + \delta_2} \overline{\text{score}}_f^\mu(T_1) + \frac{\delta_1}{\delta_1 + \delta_2} \overline{\text{score}}_f^\mu(T_2) = \frac{\delta_2(\gamma - \delta_1) + \delta_1(\gamma + \delta_2)}{\delta_1 + \delta_2} = \gamma. \quad \square$$

Acknowledgements

This research is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC), DGECR-2019-00027 and RGPIN-2019-04804.¹

References

- [Ain12] George Ainslie. Pure hyperbolic discount curves predict “eyes open” self-control. *Theory and Decision* (2012) (p. 7).
- [BB19] Eric Blais and Joshua Brody. Optimal Separation and Strong Direct Sum for Randomized Query Complexity. *Proceedings of the 34th Conference on Computational Complexity (CCC)*. 2019. DOI: [10.4230/LIPICS.CCC.2019.29](https://doi.org/10.4230/LIPICS.CCC.2019.29). arXiv: [1908.01020](https://arxiv.org/abs/1908.01020) (pp. 3, 4, 37).
- [BB20] Shalev Ben-David and Eric Blais. A Tight Composition Theorem for the Randomized Query Complexity of Partial Functions. *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 2020. DOI: [10.1109/focs46700.2020.00031](https://doi.org/10.1109/focs46700.2020.00031). arXiv: [2002.10809](https://arxiv.org/abs/2002.10809) (p. 12).
- [BB23] Shalev Ben-David and Eric Blais. A New Minimax Theorem for Randomized Algorithms. *Journal of the ACM* (6 2023). Previous version in FOCS 2020. DOI: [10.1145/3626514](https://doi.org/10.1145/3626514). arXiv: [2002.10802](https://arxiv.org/abs/2002.10802) (pp. 7, 11, 12).
- [BBGM22] Shalev Ben-David, Eric Blais, Mika Goos, and Gilbert Maystre. Randomised Composition and Small-Bias Minimax. *Proceedings of the 63rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 2022. DOI: [10.1109/focs54457.2022.00065](https://doi.org/10.1109/focs54457.2022.00065). arXiv: [2208.12896](https://arxiv.org/abs/2208.12896) (p. 7).
- [BBKW14] Amos Beimel, Sebastian Ben Daniel, Eyal Kushilevitz, and Enav Weinreb. Choosing, Agreeing, and Eliminating in Communication Complexity. *Computational Complexity* (2014). DOI: [10.1007/S00037-013-0075-7](https://doi.org/10.1007/S00037-013-0075-7) (p. 5).
- [BDG+20] Andrew Basilakis, Andrew Drucker, Mika Göös, Lunjia Hu, Weiyun Ma, and Li-Yang Tan. The Power of Many Samples in Query Complexity. *Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (ICALP)*. 2020. DOI: [10.4230/LIPIcs.ICALP.2020.9](https://doi.org/10.4230/LIPIcs.ICALP.2020.9). arXiv: [2002.10654](https://arxiv.org/abs/2002.10654) (p. 5).
- [BKST24] Guy Blanc, Caleb Koch, Carmen Strassle, and Li-Yang Tan. A Strong Direct Sum Theorem for Distributional Query Complexity. *39th Computational Complexity Conference (CCC 2024)*. 2024. DOI: [10.4230/LIPIcs.CCC.2024.16](https://doi.org/10.4230/LIPIcs.CCC.2024.16). arXiv: [2405.16340](https://arxiv.org/abs/2405.16340) (pp. 15, 31).
- [Dru12] Andrew Drucker. Improved direct product theorems for randomized query complexity. *Computational Complexity* (2012). Previous version in CCC 2011. DOI: [10.1007/s00037-012-0043-7](https://doi.org/10.1007/s00037-012-0043-7). arXiv: [1005.0644](https://arxiv.org/abs/1005.0644) (pp. 3–6, 10, 13).
- [IK10] Russell Impagliazzo and Valentine Kabanets. Constructive Proofs of Concentration Bounds. *Proceedings of the 14th International Conference on Randomization and Computation (RANDOM)*. 2010 (p. 5).

¹Cette recherche a été financée par le Conseil de recherches en sciences naturelles et en génie du Canada (CRSNG), DGEGR-2019-00027 et RGPIN-2019-04804.

- [IRW94] R. Impagliazzo, R. Raz, and A. Wigderson. A direct product theorem. *Proceedings of IEEE 9th Annual Conference on Structure in Complexity Theory*. 1994. DOI: [10.1109/SCT.1994.315814](#) (p. 3).
- [JKS10] Rahul Jain, Hartmut Klauck, and Miklos Santha. Optimal direct sum results for deterministic and randomized decision tree complexity. *Information Processing Letters* (2010). DOI: [10.1016/j.ipl.2010.07.020](#). arXiv: [1004.0105](#) (p. 3).
- [KŠW07] Hartmut Klauck, Robert Špalek, and Ronald de Wolf. Quantum and Classical Strong Direct Product Theorems and Optimal Time-Space Tradeoffs. *SIAM Journal on Computing* (2007). DOI: [10.1137/05063235X](#) (p. 3).
- [MT18] Or Meir and Avishay Tal. The choice and agreement problems of a random function. *Information Processing Letters* (2018). DOI: [10.1016/J.IPL.2017.12.007](#) (p. 5).
- [NRS98] Noam Nisan, Steven Rudich, and Michael Saks. Products and Help Bits in Decision Trees. *SIAM Journal on Computing* (1998). DOI: [10.1137/S0097539795282444](#) (p. 3).
- [Sch90] Robert Schapire. The strength of weak learnability. *Machine learning* (1990) (pp. 10, 15, 31).
- [Sha04] Ronen Shaltiel. Towards proving strong direct product theorems. *Computational Complexity* (2004). Previous version in CCC 2001. DOI: [10.1007/s00037-003-0175-x](#). ECCC: [2001/009](#) (pp. 3, 15, 31).
- [Ung09] Falk Unger. A Probabilistic Inequality with Applications to Threshold Direct-Product Theorems. *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 2009. DOI: [10.1109/FOCS.2009.62](#) (p. 5).
- [Ver98] Nikolai K. Vereshchagin. Randomized Boolean decision trees: Several remarks. *Theoretical Computer Science* (1998). DOI: [10.1016/S0304-3975\(98\)00071-1](#) (pp. 4, 37).