

---

# On Robust Streaming for Learning with Experts: Algorithms and Lower Bounds

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 In the online learning with experts problem, an algorithm makes predictions about  
2 an outcome on each of  $T$  days, given a set of  $n$  experts who make predictions on  
3 each day. The algorithm is given feedback on the outcomes of each day, including  
4 the cost of its prediction and the cost of the expert predictions, and the goal is  
5 to make a prediction with the minimum cost, compared to the best expert in  
6 hindsight. However, often the predictions made by experts or algorithms at some  
7 time influence future outcomes, so that the input is adaptively generated.

8 In this paper, we study robust algorithms for the experts problem under memory  
9 constraints. We first give a randomized algorithm that is robust to adaptive inputs  
10 that uses  $\tilde{O}\left(\frac{n}{R\sqrt{T}}\right)$  space for  $M = O\left(\frac{R^2T}{\log^2 n}\right)$ , thereby showing a smooth space-  
11 regret trade-off. We then show a space lower bound of  $\tilde{\Omega}\left(\frac{nM}{RT}\right)$  for any randomized  
12 algorithm that achieves regret  $R$  with probability  $1 - 2^{-\Omega(T)}$ , when the best expert  
13 makes  $M$  mistakes. Such an algorithm is useful for adaptive inputs, as the failure  
14 probability is low enough to union bound over all computation paths. Our result  
15 implies that the natural deterministic algorithm, which iterates through pools of  
16 experts until each expert in the pool has erred, is optimal up to polylogarithmic  
17 factors. Finally, we empirically demonstrate the benefit of using robust procedures  
18 against a white-box adversary that has access to the internal state of the algorithm.

## 19 1 Introduction

20 *Online learning with experts* is a fundamental problem in sequential prediction. On each of  $T$  days,  
21 an algorithm must make a prediction about an outcome, given a set of  $n$  experts who make predictions  
22 on the outcome. The algorithm is then given feedback on the cost of its prediction and on the expert  
23 predictions for the current day. In the *discrete prediction with experts* problem, the set of possible  
24 predictions is restricted to a finite set, and the cost is 0 if the prediction is correct, and 1 otherwise.  
25 More generally, we assume the costs are restricted to be in a range  $[0, \rho]$  for some fixed parameter  
26  $\rho > 0$ , with lower costs indicating better performance. This process continues for the  $T$  days, after  
27 which the performance (total cost) of the algorithm is compared to the performance (total cost) of  
28 the best performing expert. In particular, the goal for the online learning with experts problem is to  
29 minimize the regret, which is the amortized difference between the total cost of the algorithm and the  
30 total cost of the best performing expert, i.e., the expert that incurs the least overall cost.

31 A well-known folklore algorithm for handling the discrete prediction with experts problem is the  
32 weighted majority algorithm [42]. The deterministic variant of the weighted majority algorithm  
33 simply initializes “weights” for all experts to 1, down-weights any incorrect expert on a given day,  
34 and selects the prediction supported by the largest weight of experts. The algorithm solves the  
35 discrete prediction with experts problem with  $O(M + \log n)$  total mistakes, where  $M$  is the number

36 of mistakes made by the best expert, thus achieving regret  $O(M + \log n)$ . More generally, a large  
 37 body of literature has studied optimizations to the weighted majority algorithm, such as a randomized  
 38 variant where the probability of the algorithm selecting each prediction is proportional to the sum of  
 39 the weights of the experts supporting the prediction. The randomized weighted majority algorithm  
 40 achieves regret  $O(\sqrt{\log n/T})$  [42], which has been shown to be information-theoretically optimal,  
 41 up to a constant. There have subsequently been many follow-ups to the weighted and randomized  
 42 weighted majority algorithms that achieve similar regret bounds, but improve in other areas. For  
 43 example, on a variety of structured problems, such as online shortest paths, follow the perturbed  
 44 leader [38] achieves the same regret bound as randomized weighted majority but uses less runtime on  
 45 each day. In addition, the multiplicative weights algorithm achieves the optimal  $\sqrt{\ln n/(2T)}$  regret,  
 46 with a tight leading constant [33]. However, these classic algorithms use a framework that maintains  
 47 the cumulative cost of each expert, which requires the algorithm to store  $\Omega(n)$  bits of information.

48 **Memory bounds.** Recently, [49] considered the online learning with experts problem when memory  
 49 is a premium for the algorithm. On the hardness side, they showed that any algorithm achieving  
 50 a target average regret  $R$  requires  $\Omega(\frac{n}{R^2T})$  space, which implies that any algorithm achieving the  
 51 information-theoretic  $O(\sqrt{\log n/T})$  regret must use near-linear space. On the other hand, when the  
 52 number of mistakes  $M$  made by the best expert is small, i.e.,  $M = O(R^2T)$ , [49] gave a randomized  
 53 algorithm that uses  $\tilde{O}(\frac{n}{RT})$  space for arbitrary-order streams, thus showing that the hardness of their  
 54 lower bound originates from a setting where the best expert makes a large number of mistakes.

55 Subsequently, [47] considered the online learning with experts problem when the algorithm is limited  
 56 to memory sublinear in  $n$ . They introduced a general framework that achieves  $o(T)$  regret using  $o(n)$   
 57 memory, with a trade-off parameter between space and regret that obtains  $O_n(T^{4/5})$  regret with  
 58  $O(\sqrt{n})$  space and  $O_n(T^{0.67})$  regret with  $O(n^{0.99})$  space.

59 **Adaptive inputs.** Up to now, the discussion has focused on an oblivious setting, where the input to  
 60 the algorithm may be worst-case, but is chosen independently of the algorithm and its outputs. The  
 61 online learning with experts problem is often considered in the adaptive setting, where the input to  
 62 the algorithm is allowed to depend on previous outputs by the algorithm, e.g., in financial markets,  
 63 future stock quotes can depend on previous investment choices. Formally, we define the adaptive  
 64 setting as a two-player game between an algorithm  $\mathcal{D}$  and an adversary  $\mathcal{A}$  that adaptively creates the  
 65 input stream to  $\mathcal{D}$ . The game then proceeds in days and on the  $t$ -th day:

- 66 (1) The adversary  $\mathcal{A}$  chooses the outputs of all experts on day  $t$  as well as the outcome of day  $t$ ,  
 67 depending on all previous stream updates and all previous outputs from the algorithm  $\mathcal{D}$ .
- 68 (2) The outputs (i.e., predictions) of all experts are simultaneously given to the algorithm  $\mathcal{D}$ ,  
 69 which updates its data structures, acquires a fresh batch  $R_t$  of random bits, and outputs a  
 70 predicted outcome for day  $t$ .
- 71 (3) The outcome of day  $t$  is revealed to  $\mathcal{D}$ , while the predicted outcome for day  $t$  by  $\mathcal{D}$  is  
 72 revealed to the adversary  $\mathcal{A}$ .

73 The goal of  $\mathcal{A}$  is to induce  $\mathcal{D}$  to make as many incorrect predictions as possible throughout the stream.  
 74 It is clear that any deterministic algorithm for the online learning with experts problem will maintain  
 75 the same guarantees in the adaptive model. Unfortunately, both the algorithms of [49] and [47] are  
 76 randomized procedures that rely on iteratively sampling “pools” of experts, which can potentially be  
 77 exploited by an adaptive adversary who learns the experts sampled in each pool. Interestingly, both  
 78 the randomized weighted majority algorithm [42] and the multiplicative weights algorithm [33] are  
 79 known to be robust to adaptive inputs.

## 80 1.1 Our Contributions

81 In this paper, we study the capabilities and limits of sublinear space algorithms for the online learning  
 82 with experts problem on adaptive inputs.

83 **Robust algorithms.** Towards adaptive robustness, it is natural to study deterministic algorithms,  
 84 since they retain the same guarantee under adaptive adversaries. As a warm-up, we first provide

85 a simple deterministic algorithm that uses space  $\tilde{O}\left(\frac{nM}{RT}\right)$ . Consider an algorithm that iteratively  
 86 selects the next pool of  $k = \tilde{O}\left(\frac{nM}{RT}\right)$  experts and running the deterministic majority algorithm on the  
 87 experts in the pool, while removing any incorrect experts from the pool until the pool is completely  
 88 depleted, at which point the next pool of  $\tilde{O}\left(\frac{nM}{RT}\right)$  experts is selected. The main intuition is that each  
 89 pool can incur at most  $O(\log n)$  mistakes before it is depleted and the best expert can only make  
 90  $M$  mistakes. By the time the pool has cycled through  $nM$  experts, i.e.,  $M$  times for each of the  $n$   
 91 experts, then the best expert no longer makes any mistakes and will be retained by the pool. Thus, the  
 92 total number of mistakes made by the algorithm is  $\frac{nM}{k} \cdot O(\log n)$ . On the other hand, for a target  
 93 average regret  $R$ , the mistake bound of the algorithm is required to be at most  $M + RT$ , so it suffices  
 94 to set  $k = \tilde{O}\left(\frac{nM}{RT}\right)$  to achieve regret  $R$ . Since the algorithm runs deterministic majority on a pool of  
 95  $k = \tilde{O}\left(\frac{nM}{RT}\right)$  experts, then this algorithm uses  $\tilde{O}\left(\frac{nM}{RT}\right)$  space. Formally, we show:

96 **Theorem 1.1** (Simple deterministic algorithm; see Section 3.1). *Suppose the best expert makes  $M$*   
 97 *mistakes and let  $R \geq \frac{4M \log n}{T}$ . There exists a deterministic algorithm (Algorithm 2) that uses space*  
 98  *$\tilde{O}\left(\frac{nM}{RT}\right)$  and achieves an average regret of  $R$ .*

99 The algorithm is simple, computationally efficient, and easy to implement. However, the drawback is  
 100 that for  $M = \Omega(RT)$ , the algorithm requires space near-linear in the number of experts  $n$ , which is  
 101 undesirable when  $n$  is large. To address this issue, we complement our deterministic algorithm with  
 102 a randomized algorithm that is robust to adaptive inputs and allows for a different memory-regret  
 103 trade-off:

104 **Theorem 1.2** (Robust randomized algorithm). *Let  $R > \frac{64 \log^2 n}{T}$ , and suppose the best expert makes*  
 105 *at most  $M \leq \frac{R^2 T}{128 \log^2 n}$  mistakes. Then there exists an algorithm for the discrete prediction with*  
 106 *experts problem that uses  $\tilde{O}\left(\frac{n}{R\sqrt{T}}\right)$  space and achieves regret at most  $R$ , with high probability.*

107 This gives a trade-off between the space and regret, almost all the way to the information-theoretic  
 108 limit of  $R = O_n\left(\sqrt{1/T}\right)$  for general worst-case input. However, it incurs a multiplicative space  
 109 overhead of  $\tilde{O}(\sqrt{T})$  compared to the optimal algorithms for oblivious input. Thus we believe the  
 110 complete characterization of the space complexity of the discrete prediction with experts problem  
 111 with adaptive input is a natural open question resulting from our work.

112 **Tight memory bounds for robust algorithms.** It is natural to ask whether there exist robust  
 113 algorithms that are more space-efficient than the straightforward deterministic approach. For example,  
 114 [12] showed that any oblivious randomized algorithm with failure probability  $2^{-\Omega(nT)}$  will be  
 115 robust against adaptive outputs in the discrete prediction with experts problem, so a reasonable  
 116 approach would be to boost the success probability of existing oblivious algorithms to  $1 - 2^{-\Omega(nT)}$ .  
 117 Unfortunately, we show this cannot work:

118 **Theorem 1.3** (Memory lower bound for high-probability algorithms). *For  $n = o(2^T)$ , any random-*  
 119 *ized algorithm algorithm that achieves  $R$  regret with probability at least  $1 - 2^{-\Omega(T)}$  for the discrete*  
 120 *prediction with experts problem must use  $\Omega\left(\frac{nM}{RT}\right)$  space when the best expert makes  $M$  mistakes.*

121 In particular, Theorem 1.3 shows that any deterministic algorithm must use  $\Omega\left(\frac{nM}{RT}\right)$  space, which  
 122 taken together with the deterministic procedure above, resolves the deterministic streaming complexity  
 123 of online learning with experts. We emphasize that Theorem 1.3 also shows that using the strategy of  
 124 high-probability randomized algorithms to guarantee robustness against adaptive input does not work  
 125 any better than a deterministic algorithm.

126 At a conceptual level, our lower bound in Theorem 1.3 shows that surprisingly, the number  $M$  of the  
 127 mistakes made by the best expert is an intrinsic parameter that governs the abilities and limitations  
 128 of robust algorithms in this model. Thus, even though  $M$  is not a parameter that may naturally be  
 129 ascertained in practice, it nevertheless completely characterizes the complexity of the problem. On  
 130 the other hand, for algorithmic purposes, it suffices to acquire a constant-factor approximation to  $M$   
 131 as an input to the algorithm.

132 Another reason Theorem 1.3 is somewhat surprising is because as the number of mistakes  $M$  made  
 133 by the best expert increases, then the algorithm is also permitted to make more mistakes and in some

134 sense, the problem seems “easier”. However, [Theorem 1.3](#) shows this intuition is not true—the  
 135 problem actually becomes more difficult as  $M$  increases.

136 Moreover, we give an alternative proof in the regime when  $M = \Omega(T)$ . The proof differs from the  
 137 proof of [Theorem 1.3](#). Instead, it leverages the communication complexity of a new set disjointness  
 138 problem, recently proposed by [39]. The statement is technically weaker [Theorem 1.3](#) and appears in  
 139 the appendix; see [Appendix E](#).

140 **Empirical evaluations.** Finally, we conduct experimental evaluations in [Section 5](#) by comparing  
 141 the natural deterministic algorithm to the randomized algorithm of [49] against a white-box adversary  
 142 who has access to the internal state of the algorithm, including any experts sampled and maintained  
 143 by the algorithm. The deterministic algorithm iteratively selects pools of  $k = \tilde{O}\left(\frac{nM}{RT}\right)$  experts,  
 144 discarding any expert that has erred, and refreshing the pool with the next batch of  $k$  experts once the  
 145 pool is emptied. The randomized algorithm similarly discards erroneous experts from a pool of  $k$   
 146 experts, but it repeatedly samples pools of  $k$  experts rather than selecting the next pool of  $k$  experts.  
 147 On average across the multiple trials for each setting, the randomized algorithm made several times  
 148 more mistakes than the deterministic algorithm, ranging from 1.98x times more mistakes to 3.29x  
 149 times more mistakes than the deterministic algorithm, thus demonstrating the importance of robust  
 150 algorithms against adversarial inputs.

## 151 1.2 Related Work

152 **The experts problem and memory bounds.** The experts problem has been extensively studied [17],  
 153 both in the discrete decision setting [42] and in the setting where costs are determined by various  
 154 loss functions [35, 52–55]. Hence, the experts problem can be applied to many different applications,  
 155 such as portfolio optimization [24, 23], ensemble boosting [32], and forecasting [37]. Given certain  
 156 assumptions on the expert, such as assuming the experts are decisions trees [36, 50], threshold  
 157 functions [43], or have nice linear structures [38], additional optimizations have been made to  
 158 improve the algorithmic runtimes for the experts problem and more generally, existing work has  
 159 largely ignored optimizing for memory constraints in favor of focusing on time complexity or regret  
 160 guarantees, thus frequently using  $\Omega(n)$  memory to track the performance of each expert.

161 Recently, [49] introduced the study of memory-regret trade-offs for the experts problem. For  $n \gg T$ ,  
 162 [49] showed that the space complexity of the problem is  $\Theta\left(\frac{n}{R^2T}\right)$  in the random-order streams,  
 163 but also gave a randomized algorithm that uses  $\tilde{O}\left(\frac{n}{RT}\right)$  space for arbitrary-order streams when the  
 164 number of mistakes  $M$  made by the best expert is “small”. Subsequently, [47] considered the online  
 165 learning with experts problem for  $T \gg n$ , introducing a general space-regret trade-off framework  
 166 that achieves  $o(T)$  regret using  $o(n)$  memory, including  $O_n(T^{4/5})$  regret with  $O(\sqrt{n})$  space and  
 167  $O_n(T^{0.67})$  regret with  $O(n^{0.99})$  space.

168 **Concurrent and independent work.** Concurrent to our work, [46] considered a variant of the  
 169 problem where at each time, the algorithm selects an expert instead of a prediction. They then  
 170 introduce an algorithm robust against an adaptive adversary who observes the specific expert chosen  
 171 by the algorithm at each time, as well as lower bounds for any algorithm robust to such an adversary.

172 One way to ensure adversarial robustness is through deterministic algorithms. On that end, we  
 173 achieve stronger lower bounds for deterministic algorithms, showing that there must be a dependency  
 174 on the number  $M$  of mistakes made by the best expert, i.e., any deterministic algorithm achieving  
 175 amortized regret  $R$  must use  $\tilde{\Omega}\left(\frac{nM}{RT}\right)$  space. In fact, when the number of mistakes  $M$  made by the  
 176 best expert is sufficiently small, i.e.,  $M = O\left(\frac{R^2T}{\log^2 n}\right)$  for amortized regret  $R$ , we give a randomized  
 177 upper bound that uses *less* space than this lower bound. By comparison, the lower bound of [46]  
 178 shows that any algorithm achieving  $R$  amortized regret must use  $\tilde{\Omega}\left(\sqrt{\frac{n}{R}}\right)$  space, though their lower  
 179 bound also applies to randomized algorithms.

180 Due to the difference in setting, our algorithmic techniques are quite different from those of [46]. We  
 181 use a recent idea of [34, 4, 10] to hide the internal randomness of our algorithm from the adversary  
 182 whereas [46] rotates between groups of experts to prevent an adversary from inducing high regret by  
 183 making a specific expert bad immediately after it is selected.

184 **2 Preliminaries**

185 For any  $t \leq n$  and vector  $(X_1, X_2, \dots, X_n)$ , we let  $X_{<t}$  denote  $(X_1, \dots, X_{t-1})$ ,  $X_{\leq t} =$   
 186  $(X_1, \dots, X_t)$ , and  $X_{-t} = (X_1, \dots, X_{t-1}, X_{t+1}, \dots, X_n)$ . Also,  $X_{>t}$  and  $X_{\geq t}$  are defined simi-  
 187 larly. Let  $e_i$  denote the  $i$ th standard basis vector, and for any  $S$ ,  $e_S$  the vector that has a 1 at index  
 188  $i \in S$  and 0 everywhere else. For a random variable  $X$ , let  $H(X)$  denote its entropy.

189 We write  $[n]$  for an integer  $n > 0$  to denote the set  $\{1, \dots, n\}$ . We write  $\text{poly}(n)$  to denote a fixed  
 190 polynomial in  $n$  and if an event occurs with probability at least  $1 - \frac{1}{\text{poly}(n)}$ , we say the event occurs  
 191 with high probability. We give additional technical preliminaries in [Appendix B](#).

192 **Formal problem statement.** In the online learning with experts problem, there are  $n$  experts that  
 193 each make predictions on each of  $T$  days. The prediction are in  $\{0, 1\}$ . An algorithm uses the experts  
 194 to output a prediction for each day  $t \in [T]$ . The actual outcome of the day  $t$  is then revealed, at which  
 195 point the algorithm is penalized with a cost that is 0 if the prediction is correct, and 1 otherwise.

196 This process continues for the  $T$  days. At the end, suppose that the best expert has incurred cost  
 197  $M$ , while the algorithm has incurred  $C$ . Then the performance of the algorithm is measured by the  
 198 (average) regret  $R = \max\left(\frac{C-M}{T}, 0\right)$ .

199 **Differential privacy.** We use tools from differential privacy.

200 **Definition 2.1** (Differential privacy, [30]). *Given a privacy parameter  $\varepsilon > 0$  and a failure parameter*  
 201  *$\delta \in (0, 1)$ , a randomized algorithm  $\mathcal{A} : \mathcal{X}^* \rightarrow \mathcal{Y}$  is  $(\varepsilon, \delta)$ -differentially private if, for every pair of*  
 202 *neighboring streams  $S$  and  $S'$  and for all  $E \subseteq \mathcal{Y}$ ,*

$$\Pr[\mathcal{A}(S) \in E] \leq e^\varepsilon \cdot \Pr[\mathcal{A}(S') \in E] + \delta.$$

203 **Theorem 2.2** (Private median, e.g., [34]). *Given a database  $\mathcal{D} \in X^*$ , a privacy parameter  $\varepsilon > 0$*   
 204 *and a failure parameter  $\delta \in (0, 1)$ , there exists an  $(\varepsilon, 0)$ -differentially private algorithm PRIVMED*  
 205 *that outputs an element  $x \in X$  such that with probability at least  $1 - \delta$ , there are at least  $\frac{|S|}{2} - m$*   
 206 *elements in  $S$  that are at least  $x$ , and at least  $\frac{|S|}{2} - m$  elements in  $S$  in  $S$  that are at most  $x$ , for*  
 207  *$m = O\left(\frac{1}{\varepsilon} \log \frac{|X|}{\delta}\right)$ .*

208 **Theorem 2.3** (Advanced composition, e.g., [31]). *Let  $\varepsilon, \delta' \in (0, 1]$  and let  $\delta \in [0, 1]$ . Any mecha-*  
 209 *nism that permits  $k$  adaptive interactions with mechanisms that preserve  $(\varepsilon, \delta)$ -differential privacy*  
 210 *guarantees  $(\varepsilon', k\delta + \delta')$ -differential privacy, where  $\varepsilon' = \sqrt{2k \ln \frac{1}{\delta}} \cdot \varepsilon + 2k\varepsilon^2$ .*

211 **Theorem 2.4** (Generalization of DP, e.g., [29, 9]). *Let  $\varepsilon \in (0, 1/3)$ ,  $\delta \in (0, \varepsilon/4)$ , and  $n \geq \frac{1}{\varepsilon^2} \log \frac{2\varepsilon}{\delta}$ .*  
 212 *Suppose  $\mathcal{A} : X^n \rightarrow 2^X$  is an  $(\varepsilon, \delta)$ -differentially private algorithm that curates a database of size*  
 213  *$n$  and produces a function  $h : X \rightarrow \{0, 1\}$ . Suppose  $\mathcal{D}$  is a distribution over  $X$  and  $S$  is a set of  $n$*   
 214 *elements drawn independently and identically distributed from  $\mathcal{D}$ . Then*

$$\Pr_{S \sim \mathcal{D}, h \leftarrow \mathcal{A}(S)} \left[ \left| \frac{1}{|S|} \sum_{x \in S} h(x) - \mathbb{E}_{x \sim \mathcal{D}} [h(x)] \right| \geq 10\varepsilon \right] < \frac{\delta}{\varepsilon}.$$

215 **3 Algorithms Against Adaptive Adversaries**

216 In this section, we show that there exists algorithms for the discrete prediction with experts problem  
 217 that is robust to adaptive outputs.

218 **3.1 A Near-Optimal Deterministic Algorithm**

219 We first present a simple deterministic algorithm for arbitrary-order streams. The algorithm repeatedly  
 220 selects pools of the next  $k = \tilde{O}\left(\frac{nM}{RT}\right)$  experts. While the pool is non-empty, the algorithm runs the  
 221 deterministic majority algorithm on the algorithm and removes any incorrect experts from the pool.  
 222 Once the pool is empty, the next  $\tilde{O}\left(\frac{nM}{RT}\right)$  experts are added to the pool, possibly cycling through all  
 223  $n$  experts multiple times if necessary, where an expert can be added to the pool again even if it has  
 224 been previously deleted from the pool. We give the formal algorithm and analysis in [Appendix C](#).

225 **Theorem 3.1** (Deterministic algorithm). *Among  $n$  experts in a stream of length  $T$ , suppose the best*  
 226 *expert makes  $M$  mistakes and let  $R \geq \frac{4M \log n}{T}$ . There exists a deterministic algorithm (Algorithm 2)*  
 227 *that uses space  $\tilde{O}\left(\frac{nM}{RT}\right)$  and achieves an average regret of  $R$ .*

228 In light of lower bound [Theorem 1.3](#), it is evident that [Theorem 3.1](#) is nearly optimal, up to polyloga-  
 229 rithmic factors, for deterministic algorithms, which are automatically adversarially robust. On the  
 230 other hand, it does not seem necessary that any adversarially robust algorithm must be deterministic.  
 231 Indeed, we now give a randomized adversarially robust algorithm with better space guarantees.

### 232 3.2 A Randomized Robust Streaming Algorithm

233 We first recall the following randomized algorithm for arbitrary-order streams with oblivious input,  
 234 i.e., non-adaptive input:

235 **Lemma 3.2** (Algorithm for oblivious inputs; [49]). *Let  $R > \sqrt{\frac{128 \log^2 n}{T}}$ , and suppose the best expert*  
 236 *makes at most  $M \leq \frac{R^2 T}{1280 \log^2 n}$  mistakes. Then there exists an algorithm DISCPRED for the discrete*  
 237 *prediction with experts problem that uses  $\tilde{O}\left(\frac{n}{RT}\right)$  space and achieves regret at most  $R$ , with high*  
 238 *probability, i.e., probability at least  $1 - \frac{1}{\text{poly}(n, T)}$ .*

239 The algorithm of [Lemma 3.2](#) for constant probability proceeds by sampling pools of  $k = \tilde{O}\left(\frac{n}{RT}\right)$   
 240 experts and running majority vote on the pool, while iteratively deleting poorly performing experts  
 241 until no experts remain in the pool, at which a new pool of  $k$  experts is randomly sampled. The  
 242 main intuition is that either the pool of experts will perform well and achieve low regret, or the pool  
 243 will be continuously re-sampled until the best expert is sampled multiple times, after which point  
 244 it will not be deleted from the pool. Unfortunately, it is not evident that this algorithm is robust to  
 245 adaptive inputs because an adversary can potentially learn the experts in each sampled pool and force  
 246 the experts to make mistakes only on days in which they are sampled by the algorithm. To boost  
 247 the algorithm to high probability of success, we take the deterministic majority vote of  $O(\log n)$   
 248 independent instances of the algorithm with constant success probability.

249 Towards adaptive robustness, we use differential privacy to hide the internal randomness of the  
 250 algorithm, and in particular, the identity of the experts that are sampled by each pool. We first  
 251 run  $\tilde{O}(\sqrt{T})$  copies of the algorithm and then output the private median of the  $\tilde{O}(\sqrt{T})$  copies,  
 252 guaranteeing roughly  $\left(\frac{1}{\tilde{O}(\sqrt{T})}, 0\right)$ -differential privacy because we use  $\tilde{O}(\sqrt{T})$  copies of the algorithm.  
 253 Advanced composition, i.e., [Theorem 2.3](#), then ensures  $(O(1), 1/\text{poly}(n))$ -differential privacy, so  
 254 that correctness then follows from the generalization properties of DP, i.e., [Theorem 2.4](#).

255 We give our algorithm in full in [Algorithm 1](#).

---

**Algorithm 1** Randomized, robust streaming algorithm for the experts problem

---

**Input:** A stream of length  $T$  with  $n$  experts and a target regret  $R$

**Output:** A sequence of predictions with regret  $R$

- 1: Run  $m = O\left(\sqrt{T} \log(nT)\right)$  independent instances of DISCPRED with regret  $\frac{R}{4}$
  - 2: Run PRIVMED on the  $m$  instances with privacy parameter  $\varepsilon = O\left(\frac{1}{\sqrt{T} \log(nT)}\right)$  and failure probability  $\delta = \frac{1}{\text{poly}(n, T)}$
  - 3: At each time  $t \in [T]$ , select the output of PRIVMED
- 

256 Next, we show the correctness of our algorithm on adaptive inputs.

257 **Theorem 3.3** (Algorithm for adaptive inputs). *Let  $R > \sqrt{\frac{2048 \log^2 n}{T}}$ , and suppose the best expert*  
 258 *makes at most  $M \leq \frac{R^2 T}{1280 \log^2 n}$  mistakes. Then there exists an algorithm for the discrete prediction*  
 259 *with experts problem that uses  $\tilde{O}\left(\frac{n}{R\sqrt{T}}\right)$  space and achieves regret at most  $R$ , with probability at*  
 260 *least  $1 - \frac{1}{\text{poly}(n, T)}$ .*

261 *Proof.* Suppose we run  $m = O\left(\sqrt{T} \log(nT)\right)$  independent instances of DISCPRED with regret  $\frac{R}{4}$ .  
262 Note that for  $R > \sqrt{\frac{2048 \log^2 n}{T}}$ , we have  $\frac{R}{4} > \sqrt{\frac{128 \log^2 n}{T}}$ , which is a valid input to DISCPRED  
263 in Lemma 3.2. By Lemma 3.2, each instance succeeds on an arbitrary-order stream with probability at  
264 least  $1 - 1/\text{poly}(n, T)$ . By a union bound over the  $m$  instances, all instances succeed with probability  
265 at least  $1 - 1/\text{poly}(n, T)$ . In particular, each instance has regret at most  $R/4$ , so that the total number  
266 of mistakes by each instance is at most  $M + RT/4$ . Thus, the total number of mistakes by all  
267 instances is at most  $m(M + RT/4)$ .

268 To consider an adaptive stream, observe that PRIVMED is called with privacy parameter  
269  $O\left(1/\sqrt{T} \log(nT)\right)$  and failure probability  $1/\text{poly}(n, T)$ . By Theorem 2.3, the mechanism permits  
270  $T$  adaptive interactions and guarantees privacy  $O(1)$  with failure probability  $1/\text{poly}(n, T)$ . By The-  
271 orem 2.4, we have that with high probability, if the output of the algorithm is incorrect, then at  
272 least  $m/3$  of the instances DISCPRED are also incorrect. Since the total number of mistakes by all  
273 instances is at most  $m(M + RT/4)$ , then the total number of mistakes by the algorithm is at most  
274  $3(M + RT/4) \leq M + RT$ , since  $M \leq \frac{R^2 T}{1280 \log^2 n}$ . Hence, the algorithm achieves  $R$  regret with  
275 high probability.

276 By Lemma 3.2, each instance of DISCPRED uses  $\tilde{O}\left(\frac{n}{RT}\right)$  space. Since we use  $m =$   
277  $O\left(\sqrt{T} \log(nT)\right)$  independent instances of DISCPRED, then the total space is  $\tilde{O}\left(\frac{n}{R\sqrt{T}}\right)$ .  $\square$

## 278 4 Lower Bound for Arbitrary-Order Streams

279 In this section, we provide a space lower bound for randomized algorithms with a high probability of  
280 success. Together with Theorem 1.1, the lower bound completely characterizes the complexity of  
281 deterministic algorithms for the online learning with experts problem. We restate Theorem 1.3, give a  
282 proof sketch and defer the full analysis to Appendix D.

283 **Theorem 4.1** (Memory lower bound for high-probability algorithms). *For  $n = o(2^T)$ , any random-*  
284 *ized algorithm algorithm that achieves  $R$  regret with probability at least  $1 - 2^{-\Omega(T)}$  for the discrete*  
285 *prediction with experts problem must use  $\Omega\left(\frac{nM}{RT}\right)$  space when the best expert makes  $M$  mistakes.*

286 *Proof sketch of Theorem 4.1.* We consider the communication problem of  $\varepsilon$ -DIFFDIST. It combines  
287  $n$  instances of the distributed detection problem given by [14]. This was also used by the prior work  
288 of [49] to prove space lower bounds for expert learning in random-order stream.

289 Specifically, for fixed  $T$ , the  $\varepsilon$ -DIFFDIST problem with  $\varepsilon = \frac{M}{T}$  consists of  $T$  players, who each hold  
290  $n$  bits, indexed from 1 to  $n$ . The players must distinguish between:

- 291 (1) the NO case  $\mathcal{D}_{\text{NO}}^{(n)}$ , in which every bit for every player is drawn i.i.d. from a fair coin and
- 292 (2) the YES case  $\mathcal{D}_{\text{YES}}^{(n)}$ , in which an index  $L \in [n]$  is selected arbitrarily and the  $L$ -th bit of  
293 each player is chosen i.i.d. from a Bernoulli distribution with parameter  $\left(1 - \frac{M}{T}\right)$ , while all  
294 other bits for every player are chosen i.i.d. from a fair coin.

295 At a high level, the proof proceeds in two steps:

- 296 (1) First, we show that the  $\varepsilon$ -DIFFDIST problem can be reduced to the expert prediction problem  
297 in the streaming setting.
- 298 (2) Second, we prove a communication complexity lower bound for  $\varepsilon$ -DIFFDIST against any  
299 protocol that succeeds with probability  $1 - 2^{-\Theta(T)}$ , which includes deterministic protocols.

300 The first step is straightforward. In the reduction, each player in an instance of  $\varepsilon$ -DIFFDIST corre-  
301 sponds to a day of the expert problem. The  $n$  bit input held by each player correspond to the  $n$  expert  
302 predictions of each day. Therefore, in the NO case, each expert is correct on roughly half of the days.  
303 In the YES case, there is a single expert  $L \in [n]$  that is correct on roughly  $1/2 + \delta$  of the days (for  
304  $\delta = 1/2 - M/T$ ), while all other experts randomly guess each day. Suppose that there is a streaming

305 algorithm for the expert prediction problem with average regret  $\delta/2$ . Then roughly speaking, in the  
 306 YES case, the algorithm is correct approximately on  $1/2 + \delta/2$  of the days, while in the NO case  
 307 where every expert is randomly guessing, the algorithm is correct on less than  $1/2 + \delta/2$  of the days.  
 308 This distinguishes the YES and NO case and thus solves  $\varepsilon$ -DIFFDIST.

309 For the second step, we show that solving the  $\varepsilon$ -DIFFDIST problem with probability at least  $1 - 2^{-\Theta(T)}$   
 310 requires  $\Omega(nM)$  total communication. We give a sketch of the argument below.

311 Observe that if the input is viewed as a  $T \times n$  matrix, then  $\mathcal{D}_{\text{NO}}^{(n)}$  is a product distribution across  
 312 columns that can be written as  $\zeta^n$ , where  $\zeta$  is the distribution over a single column such that all  
 313 entries of the column are i.i.d. Bernoulli with parameter  $\frac{1}{2}$ . We view  $\mathcal{D}_{\text{NO}}^{(n)}$  as a hard distribution and  
 314 applies an information complexity analysis. By a direct sum argument, it suffices to show that the  
 315 single column problem, i.e., distinguishing between  $\mathcal{D}_{\text{NO}}^{(1)}$  and  $\mathcal{D}_{\text{YES}}^{(1)}$  (i.e., for  $n = 1$ ), requires  $\Omega(M)$   
 316 total communication.

317 Let  $(C_1, C_2, \dots, C_T)$  be a single column drawn from the hard distribution—namely, the NO case  
 318 where each player holds one i.i.d. Bernoulli with parameter  $1/2$ . Let  $A$  be a fixed protocol with  
 319 success probability at least  $1 - \exp(-\Theta(T))$ . For all  $i < T$ , let  $M_i$  denote the message sent from  
 320 player  $P_i$  to player  $P_{i+1}$  and  $M_{<i} = \{M_j : j < i\}$ . Let  $\Pi = \Pi(C_1, \dots, C_T)$  be the communication  
 321 transcript of  $A$  given the input  $(C_i)_{i=1}^T$ . A standard information complexity argument [8] implies that  
 322 the total communication is at least the *information cost*, defined as  $I(C_1, \dots, C_T; \Pi(C_1, \dots, C_T))$ ,  
 323 where  $I(X, Y)$  denotes the mutual information between random variables  $X$  and  $Y$ .

324 The key step now is to lower bound the information cost by  $\Omega(M)$ . The main ideas are the following.  
 325 For any  $i \in [T]$ , we say that  $(M_i, M_{<i})$  is *informative* for  $i$  with respect to the input  $C$  and the  
 326 transcript  $\Pi = (M_1, M_2, \dots, M_T)$  if

$$|\Pr(C_i = 0 \mid M_i, M_{<i}) - \Pr(C_i = 1 \mid M_i, M_{<i})| \geq c \quad (4.1)$$

327 for some constant  $c > 0$ . Otherwise, we say that  $M_i$  is uninformative. Intuitively, an informative  
 328 message  $M_i$  reveals sufficiently large information about  $C_i$  so that the mutual information  $I(M_i, C_i \mid$   
 329  $M_{<i})$  would be large. Let  $p_i$  be the probability that  $(M_i, M_{<i})$  is informative. Intuitively, we need  
 330 that  $\sum_i p_i$  is large, because then there would be sufficiently many informative messages, and so the  
 331 information cost is high. To formalize this approach, we claim two key lemmas. First, by [Lemma D.9](#)

$$I(\Pi; C_1, C_2, \dots, C_T) = \sum_{j=1}^T I(M_j; C_j \mid M_{<j}) \geq \Omega\left(\sum_{j=1}^T p_j\right).$$

332 Conceptually, this shows that the information cost is at least the expected number of informative  
 333 messages. Furthermore, by [Lemma D.10](#), the latter is indeed high, and in particular,  $\sum_j p_j \geq \Omega(M)$ .  
 334 Much of the technical work is dedicated to prove these lemmas. This finishes the proof since the  
 335 communication complexity is lower bounded by the information cost.  $\square$

## 336 5 Experimental Evaluations

337 In this section, we perform experimental evaluations as a simple proof-of-concept demonstrating the  
 338 importance of deterministic algorithms against adversarial input.

339 **Experimental setup.** We assume a white-box adversary with access to the internal state of the  
 340 algorithm. We evaluate the natural deterministic algorithm that iteratively selects pools of  $k =$   
 341  $\tilde{O}\left(\frac{nM}{RT}\right)$  experts, discarding any expert that has erred, and refreshing the pool with the next batch  
 342 of  $k$  experts once the pool is emptied. As a baseline, we compare to a randomized algorithm that  
 343 repeatedly samples pools of  $k = \tilde{O}\left(\frac{nM}{RT}\right)$  experts, discarding any expert that has erred, and refreshing  
 344 the pool with the next batch of  $k$  sampled experts once the pool is emptied.

345 Provided that the best expert has not yet made  $M$  mistakes, the adversary simply compels the experts  
 346 in each pool to err. Once all experts have made at least  $M$  mistakes, the adversary gives up and  
 347 permits all subsequent predictions to be correct. It can be theoretically verified that against such  
 348 an adversary, the deterministic algorithm is the optimal algorithm, in the sense that it achieves the  
 349 smallest number of errors.

350 **Experimental details.** We first evaluate our experiments on the setting  $n = 10000$ ,  $M = 20$ , and  
 351  $T = 1000$  across various values of  $R \in \{0.05, 0.1, 0.15, 0.2, 0.25, 0.3, 0.35\}$ . For each setting of  
 352  $R$ , we ran the experiment 20 times, recording the runtime and number of errors by the algorithms  
 353 in each repetition. We then computed the minimum, mean, and maximum number of errors by the  
 354 randomized algorithm across all 20 repetitions. We then repeated the experimental setup for a 10x  
 355 larger setting of  $T$ , i.e.,  $n = 10000$ ,  $M = 20$ , and  $T = 1000$ . Our experiments were performed on a  
 356 64-bit operating system using an AMD Ryzen 7 5700U CPU with 8.00 GB RAM and 8 cores with  
 357 base clock 1.80 GHz.

358 **Results.** Our experiments show that the deterministic algorithm performs significantly better than  
 359 the randomized algorithm. On average across the 20 trials for each setting, the randomized algorithm  
 360 made several times more mistakes than the deterministic algorithm, ranging from 1.98x times more  
 361 mistakes for the setting  $n = 100000$ ,  $M = 20$ ,  $T = 1000$ ,  $R = 0.05$  to 3.06x times more mistakes for  
 362 the setting  $n = 100000$ ,  $M = 10$ ,  $T = 10000$ ,  $R = 0.3$ . Even the best performance by a randomized  
 363 algorithm over all trials, which occurred at the setting  $n = 100000$ ,  $M = 20$ ,  $T = 1000$ ,  $R =$   
 364  $0.05$ , the randomized algorithm made 1.9x times more mistakes than the deterministic algorithm.  
 365 Meanwhile, the worst performance by a randomized algorithm over all trials, which occurred at  
 366 the setting  $n = 100000$ ,  $M = 10$ ,  $T = 10000$ ,  $R = 0.25$ , the randomized algorithm made 3.29x  
 367 times more mistakes than the deterministic algorithm. The average runtime was 98 seconds for  
 368 each batch of 20 experiments for the setting of  $n = 100000$ ,  $M = 20$ ,  $T = 1000$ ,  $R = 0.05$   
 369 while the average runtime was 98 seconds for each batch of 20 experiments for the setting of  
 370  $n = 100000$ ,  $M = 10$ ,  $T = 10000$ ,  $R = 0.3$  was roughly 350 seconds. See Figure 1 for a summary.

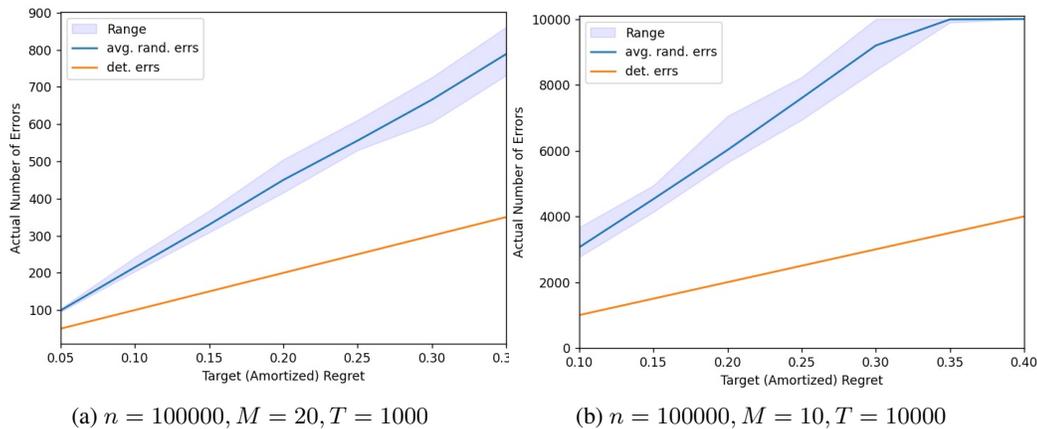


Figure 1: Comparison of errors made by deterministic algorithm and average number of errors made by randomized algorithm across 20 repetitions for each trial, across various values of input target regret  $R$ . Minimum and maximum numbers of errors by randomized algorithm across each trial are also reported.

## 371 6 Conclusion

372 In this work, we provide robust streaming algorithms for learning with experts. We provide a  
 373 deterministic algorithm parametrized by the number of mistakes made by the best expert. We also  
 374 give a randomized algorithm with a different space-regret trade-off, based on differential privacy.  
 375 We complement our algorithms with a lower bound for high-probability success algorithms. This  
 376 gives tight memory lower bound for deterministic algorithms. We then show the importance of robust  
 377 algorithmic design by empirically comparing the performance of the natural deterministic algorithm  
 378 and the state-of-the-art randomized algorithm when the inputs are adaptive.

379 We remark that our results do not rule out space-efficient robust algorithms that match the bounds of  
 380 the oblivious randomized algorithm of [49] for constant probability of success. We believe whether  
 381 or not there exists such an algorithm is a fascinating question for future work.

## References

- 382
- 383 [1] Miklós Ajtai, Vladimir Braverman, T.S. Jayram, Sandeep Silwal, Alec Sun, David P. Woodruff,  
384 and Samson Zhou. The white-box adversarial data stream model. In *Proceedings of the 41st*  
385 *ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS)*, 2022.  
386 [A](#)
- 387 [2] Raman Arora, Teodor Vanislavov Marinov, and Mehryar Mohri. Bandits with feedback graphs  
388 and switching costs. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.  
389 [A](#)
- 390 [3] Sepehr Assadi, Amit Chakrabarti, Prantar Ghosh, and Manuel Stoeckl. Coloring in graph  
391 streams via deterministic and adversarially robust algorithms. *CoRR*, abs/2212.10641, 2022. [A](#)
- 392 [4] Idan Attias, Edith Cohen, Moshe Shechner, and Uri Stemmer. A framework for adversarial  
393 streaming via differential privacy and difference estimators. In *14th Innovations in Theoretical*  
394 *Computer Science Conference, ITCS*, pages 8:1–8:19, 2023. [1.2](#), [A](#)
- 395 [5] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic  
396 multiarmed bandit problem. *SIAM J. Comput.*, 32(1):48–77, 2002. [A](#)
- 397 [6] Dmitrii Avdiukhin, Slobodan Mitrovic, Grigory Yaroslavtsev, and Samson Zhou. Adversarially  
398 robust submodular maximization under knapsack constraints. In *Proceedings of the 25th ACM*  
399 *SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD*, pages  
400 148–156, 2019. [A](#)
- 401 [7] Baruch Awerbuch and Robert D. Kleinberg. Adaptive routing with end-to-end feedback:  
402 distributed learning and geometric approaches. In *Proceedings of the 36th Annual ACM*  
403 *Symposium on Theory of Computing (STOC)*, 2004. [A](#)
- 404 [8] Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach  
405 to data stream and communication complexity. *J. Comput. Syst. Sci.*, 68(4):702–732, 2004. [4](#),  
406 [B.5](#), [B.6](#)
- 407 [9] Raef Bassily, Kobbi Nissim, Adam D. Smith, Thomas Steinke, Uri Stemmer, and Jonathan R.  
408 Ullman. Algorithmic stability for adaptive data analysis. *SIAM J. Comput.*, 50(3), 2021. [2.4](#)
- 409 [10] Amos Beimel, Haim Kaplan, Yishay Mansour, Kobbi Nissim, Thatchaphol Saranurak, and Uri  
410 Stemmer. Dynamic algorithms against an adaptive adversary: generic constructions and lower  
411 bounds. In *54th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, 2022. [1.2](#),  
412 [A](#)
- 413 [11] Omri Ben-Eliezer, Talya Eden, and Krzysztof Onak. Adversarially robust streaming via dense-  
414 sparse trade-offs. In *5th Symposium on Simplicity in Algorithms (SOSA@SODA)*, 2022. [A](#)
- 415 [12] Omri Ben-Eliezer, Rajesh Jayaram, David P. Woodruff, and Eylon Yogev. A framework for  
416 adversarially robust streaming algorithms. *SIGMOD Rec.*, 50(1):6–13, 2021. [1.1](#), [A](#)
- 417 [13] Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. Cambridge  
418 University Press, 1998. [A](#)
- 419 [14] Mark Braverman, Ankit Garg, Tengyu Ma, Huy L Nguyen, and David P Woodruff. Communica-  
420 tion lower bounds for statistical estimation problems via a distributed data processing inequality.  
421 In *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing (STOC)*,  
422 2016. [4](#)
- 423 [15] Vladimir Braverman, Avinatan Hassidim, Yossi Matias, Mariano Schain, Sandeep Silwal, and  
424 Samson Zhou. Adversarial robustness of streaming algorithms through importance sampling.  
425 In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. [A](#)
- 426 [16] Nicolò Cesa-Bianchi, Ofer Dekel, and Ohad Shamir. Online learning with switching costs and  
427 other adaptive adversaries. In *Advances in Neural Information Processing Systems (NIPS)*,  
428 2013. [A](#)

- 429 [17] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university  
430 press, 2006. 1.2, A
- 431 [18] Amit Chakrabarti, Prantar Ghosh, and Manuel Stoeckl. Adversarially robust coloring for graph  
432 streams. In *13th Innovations in Theoretical Computer Science Conference (ITCS)*, 2022. A
- 433 [19] Yeshwanth Cherapanamjeri and Jelani Nelson. On adaptive distance estimation. In *Advances  
434 in Neural Information Processing Systems 33: Annual Conference on Neural Information  
435 Processing Systems, NeurIPS*, 2020. A
- 436 [20] Yeshwanth Cherapanamjeri and Jelani Nelson. Uniform approximations for randomized  
437 hadamard transforms with applications. In *STOC '22: 54th Annual ACM SIGACT Symposium  
438 on Theory of Computing*, pages 659–671, 2022. A
- 439 [21] Yeshwanth Cherapanamjeri, Sandeep Silwal, David P. Woodruff, Fred Zhang, Qiuyi Zhang,  
440 and Samson Zhou. Robust algorithms on adaptive inputs from bounded adversaries. In *11th  
441 International Conference on Learning Representations, ICLR, Conference Track Proceedings*,  
442 2023. A
- 443 [22] Edith Cohen, Xin Lyu, Jelani Nelson, Tamás Sarlós, Moshe Shechner, and Uri Stemmer. On the  
444 robustness of counts sketch to adaptive inputs. In *International Conference on Machine Learning,  
445 ICML*, pages 4112–4140, 2022. A
- 446 [23] Thomas M Cover. Universal portfolios. *Mathematical finance*, 1(1):1–29, 1991. 1.2
- 447 [24] Thomas M. Cover and Erik Ordentlich. Universal portfolios with side information. *IEEE Trans.  
448 Inf. Theory*, 42(2):348–363, 1996. 1.2
- 449 [25] Daniela Pucci de Farias and Nimrod Megiddo. Combining expert advice in reactive environ-  
450 ments. *J. ACM*, 53(5):762–799, 2006. A
- 451 [26] Ofer Dekel, Jian Ding, Tomer Koren, and Yuval Peres. Bandits with switching costs:  $T^{2/3}$   
452 regret. In *Annual ACM Symposium on Theory of Computing (STOC)*, 2014. A
- 453 [27] Ofer Dekel, Ambuj Tewari, and Raman Arora. Online bandit learning against an adaptive  
454 adversary: from regret to policy regret. In *Proceedings of the 29th International Conference on  
455 Machine Learning (ICML)*, 2012. A
- 456 [28] Itai Dinur, Uri Stemmer, David P. Woodruff, and Samson Zhou. On differential privacy and  
457 adaptive data analysis with bounded space. *CoRR*, abs/2302.05707, 2023. A
- 458 [29] Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Leon  
459 Roth. Preserving statistical validity in adaptive data analysis. In *Proceedings of the Forty-  
460 Seventh Annual ACM on Symposium on Theory of Computing (STOC)*, 2015. 2.4
- 461 [30] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to  
462 sensitivity in private data analysis. In *Third Theory of Cryptography Conference (TCC)*, 2006.  
463 2.1
- 464 [31] Cynthia Dwork, Guy N. Rothblum, and Salil P. Vadhan. Boosting and differential privacy. In  
465 *51th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2010. 2.3
- 466 [32] Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and  
467 an application to boosting. *Journal of Computer and System Cciences*, 55(1):119–139, 1997.  
468 1.2
- 469 [33] Nick Gravin, Yuval Peres, and Balasubramanian Sivan. Tight lower bounds for multiplicative  
470 weights algorithmic families. In *44th International Colloquium on Automata, Languages, and  
471 Programming (ICALP)*, 2017. 1, 1
- 472 [34] Avinatan Hassidim, Haim Kaplan, Yishay Mansour, Yossi Matias, and Uri Stemmer. Adversari-  
473 ally robust streaming algorithms via differential privacy. In *Advances in Neural Information  
474 Processing Systems (NeurIPS)*, 2020. 1.2, 2.2, A

- 475 [35] David Haussler, Jyrki Kivinen, and Manfred K Warmuth. Tight worst-case loss bounds for  
476 predicting with expert advice. In *European Conference on Computational Learning Theory*  
477 (*EuroCOLT*), 1995. [1.2](#)
- 478 [36] David P Helmbold and Robert E Schapire. Predicting nearly as well as the best pruning of a  
479 decision tree. *Machine Learning*, 27(1):51–68, 1997. [1.2](#)
- 480 [37] Daniel A Herrmann. Prediction with expert advice applied to the problem of prediction with  
481 expert advice. *Synthese*, 200(4):1–24, 2022. [1.2](#)
- 482 [38] Adam Tauman Kalai and Santosh S. Vempala. Efficient algorithms for online decision problems.  
483 *J. Comput. Syst. Sci.*, 71(3):291–307, 2005. [1](#), [1.2](#)
- 484 [39] Akshay Kamath, Eric Price, and David P Woodruff. A simple proof of a new set disjointness  
485 with applications to data streams. In *36th Computational Complexity Conference (CCC)*, 2021.  
486 [1.1](#), [E.4](#), [E.5](#)
- 487 [40] Robert D. Kleinberg. Nearly tight bounds for the continuum-armed bandit problem. In *Advances*  
488 *in Neural Information Processing (NIPS)*, 2004. [A](#)
- 489 [41] Aryeh Kontorovich, Menachem Sadigurschi, and Uri Stemmer. Adaptive data analysis with  
490 correlated observations. In *International Conference on Machine Learning, ICML*, pages  
491 11483–11498, 2022. [A](#)
- 492 [42] Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. *Inf. Comput.*,  
493 108(2):212–261, 1994. [1](#), [1](#), [1.2](#)
- 494 [43] Wolfgang Maass and Manfred K Warmuth. Efficient learning with virtual threshold gates.  
495 *Information and Computation*, 141(1):66–83, 1998. [1.2](#)
- 496 [44] H. Brendan McMahan and Avrim Blum. Online geometric optimization in the bandit setting  
497 against an adaptive adversary. In *17th Annual Conference on Learning Theory, (COLT)*, 2004.  
498 [A](#)
- 499 [45] Neri Merhav, Erik Ordentlich, Gadiel Seroussi, and Marcelo J. Weinberger. On sequential  
500 strategies for loss functions with memory. *IEEE Trans. Inf. Theory*, 48(7):1947–1958, 2002. [A](#)
- 501 [46] Binghui Peng and Aviad Rubinfeld. Near optimal memory-regret tradeoff for online learning,  
502 2023. [1.2](#)
- 503 [47] Binghui Peng and Fred Zhang. Online prediction in sub-linear space. In *ACM-SIAM Symposium*  
504 *on Discrete Algorithms (SODA)*, 2023. [1](#), [1](#), [1.2](#), [A](#)
- 505 [48] Chloé Rouyer, Yevgeny Seldin, and Nicolò Cesa-Bianchi. An algorithm for stochastic and  
506 adversarial bandits with switching costs. In *Proceedings of the 38th International Conference*  
507 *on Machine Learning (ICML)*, 2021. [A](#)
- 508 [49] Vaidehi Srinivas, David P. Woodruff, Ziyu Xu, and Samson Zhou. Memory bounds for the  
509 experts problem. In *54th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*,  
510 2022. [1](#), [1](#), [1.1](#), [1.2](#), [3.2](#), [4](#), [6](#), [D.1](#), [D.1](#), [D.1](#), [D.3](#), [D.1](#)
- 511 [50] Eiji Takimoto, Akira Maruoka, and Vladimir Vovk. Predicting nearly as well as the best pruning  
512 of a decision tree through dynamic programming scheme. *Theor. Comput. Sci.*, 261(1):179–209,  
513 2001. [1.2](#)
- 514 [51] Flemming Topsøe. Bounds for entropy and divergence for distributions over a two-element set.  
515 *J. Ineq. Pure & Appl. Math*, 2(2):300, 2001. [B.1](#)
- 516 [52] Vladimir Vovk. Aggregating strategies. In *Proceedings of the Third Annual Workshop on*  
517 *Computational Learning Theory (COLT)*, 1990. [1.2](#)
- 518 [53] Vladimir Vovk. A game of prediction with expert advice. *J. Comput. Syst. Sci.*, 56(2):153–173,  
519 1998.

- 520 [54] Vladimir Vovk. Derandomizing stochastic prediction strategies. *Machine Learning*, 35(3):247–  
521 282, 1999.
- 522 [55] Vladimir Vovk. Defensive prediction with expert advice. In *16th International Conference on*  
523 *Algorithmic Learning Theory (ALT)*, 2005. [1.2](#)
- 524 [56] David Wajc. Rounding dynamic matchings against an adaptive adversary. In *52nd Annual ACM*  
525 *SIGACT Symposium on Theory of Computing (STOC)*, 2020. [A](#)
- 526 [57] David P. Woodruff and Samson Zhou. Tight bounds for adversarially robust streams and  
527 sliding windows via difference estimators. In *62nd IEEE Annual Symposium on Foundations of*  
528 *Computer Science (FOCS)*, 2021. [A](#)

## 529 A Additional Related Work on Adaptive Inputs

530 Motivated by non-independent inputs and adversarial attacks, adaptive inputs have recently been  
531 considered in the centralized model [19, 41, 20, 21], in the streaming model [6, 12, 34, 57, 15, 18,  
532 11, 1, 3, 22, 4, 28], and in the dynamic model [56, 10]. In particular, algorithms robust to inputs  
533 that can depend on the previous outputs by the algorithm, i.e., black-box attacks, are also robust to  
534 situations in which future inputs may be dependent on previous outputs. This is especially relevant in  
535 applications such as forecasting, in which a prediction on day  $i$  can lead to a series of actions that  
536 might impact outcomes and expert predictions on day  $i + 1$  and beyond.

537 Adaptive adversaries have received considerable attention in literature for online learning when  
538 the goal is simply to achieve the best possible regret [13, 17, 44]. Building off a line of results on  
539 multi-armed bandit problems [5, 7, 40], the work of [45] first considered the experts setting against  
540 memory-bounded adaptive adversaries, giving an algorithm with regret  $O(T^{2/3})$ . An early paper of  
541 [25] introduced a family of algorithms for adaptive inputs, but provided guarantees using concepts not  
542 quite related to the standard definitions of regret. More recent works have explored online learning  
543 with additional considerations, such as alternative quantities to optimize [27], additional switching  
544 costs [16, 26, 48], and feedback graphs [2]. The closest work to our setting is the recent result by [47]  
545 showing that no algorithm using space sublinear in  $n$  can achieve regret sublinear in  $T$  when the  
546 input is chosen by an adversary with access to the internal state of the algorithm, i.e., a white-box  
547 adversary.

## 548 B Additional Technical Preliminaries

### 549 B.1 Information Theory

550 For any  $p \in [0, 1]$ , we slightly abuse notation and let  $H(p) = -p \log_2 p - (1 - p) \log_2 (1 - p)$  be the  
551 binary entropy function. The following is a standard upper and lower bound of  $H(p)$ .

552 **Lemma B.1** (Bound on the binary entropy function; see e.g. [51]). *For  $p \in [0, 1]$ , the binary entropy  
553 function satisfies*

$$4p(1 - p) \leq H(p) \leq 2(p(1 - p))^{1/\ln 4}.$$

### 554 B.2 Communication Complexity

555 **Definition B.2** (Mutual information). *Let  $X$  and  $Y$  be a pair of random variables with joint  
556 distribution  $p(x, y)$ . Then the mutual information is defined as  $I(X; Y) := \sum_{x, y} p(x, y) \log \frac{p(x, y)}{p(x)p(y)}$ ,  
557 for marginal distributions  $p(x)$  and  $p(y)$ .*

558 In a multi-party communication problem of  $t$  players, each player is given  $x_i \in \mathcal{X}_t$ . They communi-  
559 cate according to fixed protocol to compute a function  $f : \mathcal{X}_t \times \dots \times \mathcal{X}_t \rightarrow \mathcal{Y}$ . A protocol  $\Pi$  is called  
560 a  $\delta$ -error protocol for  $f$  if there exists a function  $\Pi_{\text{out}}$  such that  $\Pr [\Pi_{\text{out}}(\Pi(x, y)) = f(x, y)] \geq 1 - \delta$ .  
561 For a (multi-party) communication problem, we denote the transcript of all communication in a  
562 protocol as  $\Pi \in \{0, 1\}^*$ . The communication cost of a protocol, as a result, is the bit length of the  
563 transcript. Let  $R_\delta(f)$  denote the minimum communication cost across all  $\delta$ -error protocols for  $f$ .

564 **Definition B.3** (Information cost). *Let  $\Pi$  be a randomized protocol that produces a random variable  
565  $\Pi(X_1, \dots, X_T)$  as a transcript on inputs  $X_1, \dots, X_T$  drawn from a distribution  $\mu$ . Then the  
566 information cost of  $\Pi$  with respect to  $\mu$  is defined as  $I(X_1, \dots, X_T; \Pi(X_1, \dots, X_T))$ .*

567 **Definition B.4** (Information complexity). *The information complexity of a function  $f$  with respect to  
568 a distribution  $\mu$  and failure probability  $\delta$  is the minimum information cost of a protocol for  $f$  with  
569 respect to  $\mu$  that fails with probability at most  $\delta$  on every input and denoted by  $\text{IC}_{\mu, \delta}(f)$ .*

570 **Lemma B.5** (Information cost decomposition lemma, Lemma 5.1 in [8]). *Let  $\mu$  be a mixture  
571 of product distributions and suppose  $\Pi$  is a protocol for inputs  $(X_1, \dots, X_T) \sim \mu^n$ . Then  
572  $I(X_1, \dots, X_T; \Pi(X_1, \dots, X_T)) \geq \sum_{i=1}^n I(X_{1,i}, \dots, X_{T,i}; \Pi(X_1, \dots, X_T))$ , where  $X_{i,j}$  denotes  
573 the  $j$ -th component of  $X_i$ .*

574 **Lemma B.6** (Information complexity lower bounds communication complexity; Proposition 4.3 [8]).  
575 *For any distribution  $\mu$  and error  $\delta$ ,  $R_\delta(f) \geq \text{IC}_{\mu, \delta}(f)$ .*

576 **C Proof of Theorem 3.1 and Formal Algorithm**

577 We give a formal description of our deterministic robust algorithm in pseudocode.

---

**Algorithm 2** Deterministic algorithm for the experts problem

---

**Input:** A stream of length  $T$  with  $n$  experts, upper bound  $M$  on the number of mistakes made by the best expert, and target regret  $R$   
**Output:** A sequence of predictions with regret  $R$

- 1:  $k \leftarrow \frac{4nM}{RT} \log n$
- 2:  $S \leftarrow \emptyset$
- 3: **while** the stream persists **do**
- 4:     **if**  $S$  is empty **then** ▷ We have cycled through all  $n$  experts once
- 5:          $S \leftarrow [n]$
- 6:     Let  $P$  be the first  $k$  indices of  $S$
- 7:      $S \leftarrow S \setminus P$
- 8:     **while**  $P \neq \emptyset$  **do**
- 9:         For each following day, choose the outcome output by the majority of the experts in  $P$
- 10:         Delete the incorrect experts on that day from  $P$

---

578 We now prove the correctness and space complexity of Algorithm 2.

579 *Proof of Theorem 3.1.* We first remark that the algorithm can make at most  $\log k \leq \log n$  mistakes  
580 over the lifespan of each pool of size  $k := \frac{2nM}{RT} \log n$  because each time the algorithm makes a  
581 mistake, at least half of the pool must be incorrect and deleted, so the size of the pool decreases by  
582 at least half with each mistake the algorithm makes. Note that  $k \leq n$  for  $R \geq \frac{4M \log n}{T}$ , so the  
583 algorithm is well-defined.

584 Since each pool  $P$  has size  $k$  and there are  $n$  experts, then there are at most  $\frac{4n}{k}$  pools before the entire  
585 set  $S$ , which is initialized to  $n$ , is depleted. Thus, there are at most  $\frac{4n}{k}$  pools to iterate through the  
586 entire set of experts. Moreover, each time the algorithm has iterated through the entire set of experts,  
587 each expert must have made at least one mistake. This is because an expert is only deleted from the  
588 pool  $P$  when it has made a mistake and since all experts have been deleted from  $P$ , then all experts  
589 have made at least one mistake.

590 Since the best expert makes at most  $M$  mistakes, then the best expert can be deleted from the pool  $P$   
591 at most  $M$  times. In other words, the algorithm can cycle through the entire set of  $n$  experts at most  
592  $M + 1$  times.

593 Hence, the total number of mistakes by the algorithm is at most

$$\frac{2n}{k} \cdot \log n \cdot (M + 1) \leq \frac{4n}{k} \cdot \log n \cdot M \leq \frac{4nRT}{4nM \log n} \cdot \log n \cdot M = RT,$$

594 so the algorithm achieves regret at most  $R$ . Since the algorithm selects a subset of  $k = \frac{4nM}{RT} \log n$   
595 experts, then the space complexity follows.  $\square$

596 **D Proofs of the Lower Bounds for Arbitrary-Order Streams**

597 In this section, we give space lower bounds for the experts problem on arbitrary-order streams. As  
598 a warm-up, we first show in Section D.1 a general space lower bound for randomized algorithms  
599 when the best expert makes a “small” number of mistakes. We then give our main lower bound result  
600 in Section D.2, showing that any deterministic algorithm achieving regret  $R$  must use space  $\Omega\left(\frac{nM}{RT}\right)$   
601 when the best expert makes  $M$  mistakes.

602 **D.1 Warm-up: Lower Bound for Accurate Best Expert**

603 In this section, we show that any randomized algorithm that achieves regret  $R$  must use  $\Omega\left(\frac{n}{RT}\right)$   
604 space, even when the best expert makes  $\Theta(RT)$  mistakes. In contrast, [49] give an  $\Omega\left(\frac{n}{R^2T}\right)$  space  
605 lower bound:

606 **Theorem D.1** (Memory lower bound; Theorem 1 of [49]). *Let  $R > 0$ ,  $p < \frac{1}{2}$  be fixed constants, i.e.,*  
607 *independent of other input parameters. Any algorithm that achieves  $R$  regret for the experts problem*  
608 *with probability at least  $1 - p$  must use at least  $\Omega\left(\frac{n}{R^2T}\right)$  space.*

609 *Furthermore, this lower bound holds even when the costs are binary, and expert predictions, as well*  
610 *as the correct answers, are constrained to be i.i.d. across the days, albeit with different distributions*  
611 *across the experts.*

612 The proof of this lower bound exploits a construction where the best expert makes  $\Theta(T)$  mistakes.  
613 Thus, it is not clear how the space complexity of the problem behaves when the best expert makes a  
614 smaller number of mistakes. In fact, [49] also give an algorithm that uses  $\tilde{O}\left(\frac{n}{RT}\right)$  space when the  
615 best expert makes  $O(RT)$  mistakes, bypassing the aforementioned lower bound.

616 We now prove that in this small mistake regime, this algorithm is tight. Towards this goal, we first  
617 define the  $\varepsilon$ -DIFFDIST problem that reduces to the experts problem. It was proposed by [49] to prove  
618 memory lower bounds for the expert problem in random order stream.

619 **Definition D.2** (The  $\varepsilon$ -DIFFDIST Problem). *We have  $T$  players, each of whom holds  $n$  bits, indexed*  
620 *from 1 to  $n$ . We must distinguish between two cases, which we refer to as “ $V = 0$ ” and “ $V = 1$ ”. Let*  
621  *$\mu_0$  be a Bernoulli distribution with parameter  $\frac{1}{2}$ , i.e., a fair coin, and let  $\mu_1$  be a Bernoulli distribution*  
622 *with parameter  $\frac{1}{2} + \varepsilon$ .*

- 623 • (NO Case, “ $V = 0$ ”) *Every index for every player is drawn i.i.d. from a fair coin, i.e.,  $\mu_0$ .*
- 624 • (YES Case, “ $V = 1$ ”) *An index  $L \in [n]$  is selected arbitrarily—the  $L$ -th bit of each player*  
625 *is chosen i.i.d. from  $\mu_1$ . All other bits for every player are chosen i.i.d. from  $\mu_0$ .*

626 Any protocol that successfully solves the  $\varepsilon$ -DIFFDIST problem with a constant probability greater  
627 than  $\frac{1}{2}$  must use at least  $\Omega\left(\frac{n}{\varepsilon^2}\right)$  communication, a result due to [49]:

628 **Lemma D.3** (Communication complexity of  $\varepsilon$ -DIFFDIST; Lemma 3 of [49]). *The communication*  
629 *complexity of solving the  $\varepsilon$ -DIFFDIST problem with a constant  $1 - p$  probability, for any  $p \in [0, 0.5)$ ,*  
630 *is  $\Omega\left(\frac{n}{\varepsilon^2}\right)$ .*

631 The proof of [Theorem D.1](#) by [49] uses  $n$  coin flips across each of the  $T$  players to form the  $n$  expert  
632 predictions over each of the  $T$  days. In the NO case, each expert will be correct on roughly  $\frac{T}{2}$  days,  
633 while in the YES case, a single expert will be correct on roughly  $\frac{T}{2} + \varepsilon T$  days, so that an algorithm  
634 with regret  $R = O(\varepsilon)$  will be able to distinguish between the two cases. There is a slight subtlety  
635 in the proof that uses a masking argument to avoid “trivial” algorithms that happen to succeed on a  
636 “lucky” input, but for the purposes of our proof in this section, the masking argument is not needed.  
637 It then follows that the total communication is  $\Omega\left(\frac{n}{R^2}\right)$  across the  $T$  players, so that any streaming  
638 algorithm must use at least  $\Omega\left(\frac{n}{R^2T}\right)$  bits of space.

639 Suppose we instead consider the  $\varepsilon$ -DIFFDIST problem over  $RT$  players, representing  $RT$  days in  
640 the experts problem. Moreover, suppose we set  $\varepsilon = \Theta(1)$  in the  $\varepsilon$ -DIFFDIST problem, so that in  
641 the NO case, each of the experts will be correct on roughly  $\frac{RT}{2}$  days, while in the YES case, a  
642 single expert will be correct on roughly  $\frac{RT}{2} + CRT$  days, for some constant  $C > 0$ . Suppose we  
643 further pad all of the experts with incorrect predictions across an additional  $T - RT$  days, so that the  
644 total number of days is  $T$ , but the number of correct expert predictions remains the same. Then an  
645 algorithm achieving regret  $O(R)$  will be able to distinguish between the two cases, so that the total  
646 communication is  $\Omega\left(\frac{n}{R}\right)$ , so that any streaming algorithm must use at least  $\Omega\left(\frac{n}{RT}\right)$  bits of space.

647 **Corollary D.4.** *Let  $R$ ,  $p < \frac{1}{2}$  be fixed constants, i.e., independent of other input parameters. Any*  
648 *algorithm that achieves  $R$  regret for the experts problem with probability at least  $1 - p$  must use*  
649 *at least  $\Omega\left(\frac{n}{RT}\right)$  space even when the best expert makes as few as  $\Theta(RT)$  mistakes. This lower*  
650 *bound holds even when the costs are binary and expert predictions, as well as the correct answer, are*  
651 *constrained to be i.i.d. across the days, albeit with different distributions across the experts.*

652 *Proof.* The claim follows from setting  $T = RT$  and  $R = \Theta(1)$  in the proof of [Theorem D.1](#). □

653 **D.2 Lower Bound for Deterministic Algorithms**

654 We now prove our main space lower bound for deterministic algorithms (Theorem 1.3). We first set  
655 up some basic notations and introduce a hard distribution.

656 Let  $T$  be any fixed positive integer. Let  $\mathcal{D}_{\text{NO}}^{(n)}$  be the distribution over matrices  $A$  with size  $T \times n$   
657 such that all entries of the matrix are i.i.d. Bernoulli with parameter  $\frac{1}{2}$ , i.e., each entry of  $A$  is 0  
658 with probability  $\frac{1}{2}$  and 1 with probability  $\frac{1}{2}$ . Let  $\mathcal{D}_{\text{YES}}^{(n)}$  be the distribution over matrices  $M$  with size  
659  $T \times n$  such that there is a randomly chosen column  $L \in [n]$ , which is i.i.d. Bernoulli with parameter  
660  $(1 - \frac{M}{T})$  and all other columns are i.i.d. Bernoulli with parameter  $\frac{1}{2}$ . Let  $\text{BIASDETECT}_n$  be the  
661 problem of detecting whether  $A$  is drawn from  $\mathcal{D}_{\text{YES}}^{(n)}$  or  $\mathcal{D}_{\text{NO}}^{(n)}$ , so that  $\text{BIASDETECT}_n$  is simply the  
662  $\varepsilon$ -DIFFDIST problem with  $\varepsilon = \frac{1}{2} - \frac{M}{T}$ .

663 Let  $\Pi$  be a communication protocol for  $\text{BIASDETECT}_n$  that is correct with probability at least  
664  $1 - \exp(-\Theta(T))$ . Since  $\mathcal{D}_{\text{NO}}^{(n)}$  is a product distribution across columns, then it can be written as  $\zeta^n$ ,  
665 where  $\zeta$  is the distribution over a single column such that all entries of the column are i.i.d. Bernoulli  
666 with parameter  $\frac{1}{2}$ . Let  $\text{BIASDETECT}_1$  denote the problem of distinguishing between  $\mathcal{D}_{\text{NO}}^{(1)}$  and  $\mathcal{D}_{\text{YES}}^{(1)}$   
667 on a single column, i.e.,  $n = 1$ . Using  $\mathcal{D}_{\text{NO}}^{(n)}$  as the hard distribution, we have the following direct  
668 sum theorem.

669 **Lemma D.5** (Direct sum for BIASDETECT). *The information complexity of  $\text{BIASDETECT}_n$  satisfies*

$$\text{IC}_{\mathcal{D}_{\text{NO}}^{(n)}, 2^{-\Theta(T)}}(\text{BIASDETECT}_n) \geq n \cdot \text{IC}_{\mathcal{D}_{\text{NO}}^{(1)}, 2^{-\Theta(T)}}(\text{BIASDETECT}_1).$$

670 *Proof.* By definition,  $\mathcal{D}_{\text{NO}}^{(n)} = \zeta^n$  is a product distribution over  $n$  columns. The lemma follows from  
671 the standard direct sum lemma of information cost (Lemma B.5).  $\square$

672 With the above direct sum theorem for  $\text{BIASDETECT}_n$ , it now suffices to provide a single-coordinate  
673 information cost lower bound against  $\text{BIASDETECT}_1$ . The proof is delayed to Section D.3.

674 **Lemma D.6** (Single-coordinate information cost lower bound). *Let  $c \in (0, 1)$  and  $\Pi$  be any protocol  
675 with error  $\delta = 2^{-\Theta(T)}$  for  $\text{BIASDETECT}_1$ . We have that the information cost of  $\Pi$  with respect to  $\zeta$   
676 is at least*

$$I(\Pi(C_1, C_2, \dots, C_T); C_1, C_2, \dots, C_T) \geq \Omega(M), \quad (\text{D.1})$$

677 where the bits  $C_i \sim \zeta$  are i.i.d. single coordinates.

678 Combining Lemma D.6 with the direct sum theorem (Lemma D.5), we immediately get the following  
679 information complexity lower bound for  $\text{BIASDETECT}_n$ :

**Theorem D.7** ( $n$ -Coordinate information complexity lower bound). *Let  $c \in (0, 1)$ . Then*

$$\text{IC}_{\mathcal{D}_{\text{NO}}^{(n)}, 2^{-\Theta(T)}}(\text{BIASDETECT}_n) = \Omega(nM).$$

680 *Proof.* This follows by applying the direct sum theorem (Lemma D.5) to the single-coordinate bound  
681 Lemma D.6.  $\square$

682 This implies that any algorithm with  $R$  regret and success rate at least  $1 - 2^{-\Theta(T)}$  requires  $\Omega(\frac{nM}{RT})$   
683 memory, where  $M$  is the mistake bound on the best expert.

684 **Theorem D.8** (Memory lower bound for expert learning). *Let  $R, M$  be fixed and independent of  
685 other input parameters. Any streaming algorithm that achieves  $R$  regret for the experts problem  
686 with probability at least  $1 - 2^{-\Theta(T)}$  must use at least  $\Omega(\frac{nM}{RT})$  space, for  $n = o(2^T)$ , where the best  
687 expert makes  $M$  mistakes.*

688 *Proof.* We now consider the problem  $\text{BIASDETECT}_n$  on a matrix of size  $RT \times n$ . Note that in the  
689 NO case, at any fixed column  $i \in [n]$ , the probability that there are more than  $\frac{3RT}{5} - \frac{M}{2}$  instances of  
690 0, for  $M \leq \frac{RT}{8}$ , is at most  $2 \exp(-c_1 RT)$ , for a sufficiently small constant  $c_1 \in (0, 1)$ . Thus, by a  
691 union bound, the probability that there exists an index  $i \in [n]$  with more than  $\frac{3RT}{4} - \frac{M}{2}$  instances of  
692 0 is at most  $2n \exp(-c_1 RT)$ .

693 Similarly in the YES case, the probability that there are fewer than  $\frac{4RT}{5} - \frac{M}{2}$  instances of 0 for a  
694 fixed  $i \in [n]$  and for  $M \leq \frac{RT}{8}$  is at most  $2 \exp(-c_2 RT)$ , for a sufficiently small constant  $c_2 \in (0, 1)$   
695 and so by a union bound, the probability that there exists an index  $i \in [n]$  with fewer than  $\frac{3RT}{4} - \frac{M}{2}$   
696 instances of 0 is at most  $2n \exp(-c_2 RT)$ . Hence, for  $n = o(2^T)$ , there exists a constant  $c \in (0, 1)$   
697 such that any algorithm that achieves total regret at most  $\frac{RT}{5}$  with probability at least  $1 - \exp(-cT)$   
698 can distinguish between the YES and NO cases with probability  $1 - \exp(-\Theta(T))$ .

699 By [Theorem D.7](#) and [Lemma B.6](#), the total communication across the  $RT$  players must be at least  
700  $\Omega(nM)$ . Therefore, any streaming algorithm that achieves average  $R$  regret for the experts problem  
701 with probability at least  $1 - 2^{-\Theta(T)}$  must use at least  $\Omega(\frac{nM}{RT})$  space.  $\square$

### 702 D.3 Proof of the Single-Coordinate Information Cost Lower Bound

703 We now show the single-coordinate lower bound of [Lemma D.6](#).

704 *Proof of Lemma D.6.* Consider a protocol that is correct with probability  $1 - 2^{-\Theta(T)}$  and let  
705  $(C_1, C_2, \dots, C_T) \sim \zeta^T$  be a single column drawn from the NO case, where each coordinate is  
706 i.i.d. Bernoulli with parameter  $1/2$ . For notational convenience, let  $\Pi = \Pi(C_1, \dots, C_T)$  denote  
707 the transcript given the input  $(C_1, C_2, \dots, C_T)$ . We consider the one-way message-passing model,  
708 where each player  $P_i$  holds the input  $C_i$ . For all  $i < T$ , let  $M_i$  denote the message sent from player  
709  $P_i$  to player  $P_{i+1}$ .

710 By the chain rule of mutual information, the information cost of the transcript, the left-side of  
711 [Equation D.1](#) that we need to bound, can be written as

$$I(\Pi; C_1, C_2, \dots, C_T) = \sum_{j=1}^T I(M_j; C_1, C_2, \dots, C_T \mid M_{<j}). \quad (\text{D.2})$$

712 By the independence of one-way communication, we have

$$I(M_j; C_1, C_2, \dots, C_T \mid M_{<j}) = I(M_j; C_j \mid M_{<j}). \quad (\text{D.3})$$

713 Combining the two equalities above, the information cost equals

$$I(\Pi; C_1, C_2, \dots, C_T) = \sum_{j=1}^T I(M_j; C_j \mid M_{<j}). \quad (\text{D.4})$$

714 We now lower bound the right-side. First, we make the following definition. For any  $i \in [T]$ , we say  
715 that  $(M_i, M_{<i})$  is *informative* for  $i$  with respect to the input  $C$  and the transcript  $\Pi = (M_1, \dots, M_T)$   
716 if

$$|\Pr(C_i = 0 \mid M_i, M_{<i}) - \Pr(C_i = 1 \mid M_i, M_{<i})| \geq c \quad (\text{D.5})$$

717 for some constant  $c > 0$ ; and uninformative otherwise. Intuitively, an informative index  $i$  with respect  
718 to  $(M_i, M_{<i})$  means that conditional on the past messages  $M_{<i}$ , the message  $M_i$  reveals much  
719 information about  $C_i$ . Hence, in this case,  $I(M_i, C_i \mid M_{<i})$  would be large. Now for all  $i \in [T]$ , let  
720  $p_i$  be the probability that  $(M_i, M_{<i})$  is informative (for  $i$  with respect to  $C$  and  $\Pi$ ).

721 Conceptually, we need to show that  $\sum_i p_i$  is large, since then there would be sufficiently many  
722 informative messages, and so the information cost in the left-side of [Equation D.4](#) is high. We  
723 formalize this idea in the following lemma.

724 **Lemma D.9.** *In the setting above, where  $c > 0$  is a constant, the information cost can be lower*  
725 *bounded by*

$$I(\Pi; C_1, C_2, \dots, C_T) = \sum_{j=1}^T I(M_j; C_j \mid M_{<j}) \geq \Omega\left(\sum_{j=1}^T p_j\right) \quad (\text{D.6})$$

726 *Proof.* We start by expanding the definition of the mutual information terms. For each  $j \in T$ , we  
727 have

$$I(M_j; C_j \mid M_{<j}) = H(C_j \mid M_{<j}) - H(C_j \mid M_j, M_{<j}) \quad (\text{D.7})$$

728 For the first term, notice that  $C_j$  and  $M_{<j}$  are independent by one-way communication. Moreover,  
 729 by definition  $C_j$  is Bernoulli with parameter  $1/2$ . Therefore,

$$H(C_j | M_{<j}) = H(C_j) = H(1/2) = 1.$$

730 For the second term,

- 731 • either  $(M_j, M_{<j})$  is informative, which holds with probability  $p_j$ , and in this case, the  
 732 conditional entropy is upper bounded by  $H(C_j | M_j, M_{<j}) \leq H(1/2 + c/2)$ ;
- 733 • or  $(M_j, M_{<j})$  is uninformative, and in this case, we trivially upper bound the conditional  
 734 entropy by  $H(C_j | M_j, M_{<j}) \leq 1$ ;

735 Putting the observations together and using Equation D.7, it follows that

$$\begin{aligned} I(M_j; C_j | M_{<j}) &= H(C_j | M_{<j}) - H(C_j | M_j, M_{<j}) \\ &\geq 1 - (p_j \cdot H(1/2 + c/2) + (1 - p_j) \cdot 1) \\ &= p_j - p_j \cdot H(1/2 + c/2) \\ &\geq p_j \left(1 - (1 - c^2)^{1/\ln 4}\right). \end{aligned}$$

736 where the last step uses the upper bound of Lemma B.1. Then we have

$$\begin{aligned} I(M_j; C_j | M_{<j}) &\geq p_j \left(1 - (1 - c^2)^{1/\ln 4}\right) \\ &\geq c^3 \cdot \Omega(p_j), \end{aligned}$$

737 where the last step follows since  $1 - (1 - x^2)^{1/\ln 4} \geq x^3/100$  for  $x \in [0, 1]$ . Summing over  
 738  $j = 1, 2, \dots, T$  in Equation D.6 finishes the proof.  $\square$

739 To prove the claimed information cost inequality Equation D.1, we show that  $\sum_i p_i = \Omega(M)$ .

740 **Lemma D.10.** *There exists a constant  $\gamma > 0$  such that*

$$\sum_{j=1}^T p_j > \gamma \cdot M.$$

741 *Proof.* Suppose by way of contradiction that  $\sum_{j=1}^T p_j = o(M)$ . Let  $A$  be a protocol that sends  
 742 (possibly random) messages  $M_1, \dots, M_T$  on a random input  $C \in \{0, 1\}^T \sim \zeta^T$  drawn uniformly  
 743 from the NO distribution, i.e., each coordinate of  $C := C_1, \dots, C_T$  is picked to be 0 with probability  
 744  $\frac{1}{2}$  and 1 with probability  $\frac{1}{2}$ . Moreover, suppose  $A$  is a protocol that distinguishes between a YES  
 745 instance and a NO instance with probability at least  $1 - \frac{e^{-cT}2^{-T}}{8}$ , for some constant  $c > 0$ .

746 Since  $p_i$  is the probability that  $M_i$  is informative, then by assumption, the expected number of  
 747 informative indices  $i$  over the messages  $M_1, \dots, M_T$  is  $f(M)$  for some  $f(M) = o(M)$ . Thus by  
 748 Markov's inequality, the probability that the number of informative indices is at most  $10f(M) =$   
 749  $o(M)$  with probability at least  $\frac{9}{10}$ . Let  $S$  be the set of the uninformative indices so that  $|S| =$   
 750  $T - 10f(M) = T - o(M)$ . Let  $C'$  be an input that agrees with  $C$  on the informative indices  $[T] \setminus S$   
 751 and is chosen arbitrarily on uninformative indices  $S$ , so that  $C'_i = C_i$  for  $i \in [T] \setminus S$ .

752 By definition, each uninformative index only changes the distribution of the output by a  $(1 \pm c)$  factor.  
 753 In particular, for  $c \in (0, 1/2)$ , the probability that the protocol  $A$  generates  $\Pi$  on input  $C'$  is at least  
 754  $(1 - c)^T \geq e^{-2cT}$  times the probability that the protocol  $A$  generates  $\Pi$  on input  $C$ . However, since  
 755  $C$  can differ from  $C'$  on  $S$ , then  $C$  can differ from  $C'$  on  $|S| = T - 10f(M) = T - o(M)$  indices.

756 Now since each coordinate of  $C$  is picked to be 0 with probability  $\frac{1}{2}$  and 1 with probability  $\frac{1}{2}$ , then  
 757 the probability that  $C$  contains more than  $T - M$  zeros is at least  $1 - T^M \cdot \frac{1}{2^T} \geq 1 - 2^{T/2}$  for  
 758 sufficiently large  $T$ . But then there exists a choice of  $C'$  that contains fewer than  $\frac{M}{2}$  zeros such that  
 759  $A$  will also output  $\Pi$  with probability at least  $\frac{e^{-cT}}{2}$ . Since  $C'$  contains fewer than  $\frac{M}{2}$ , then  $C'$  is more  
 760 likely to be generated from a YES instance and indeed a YES instance will generate  $C$  with probability

761  $2^{-T}$ . On the other hand, since  $\Pi$  corresponds to a transcript for which  $A$  will output NO, then the  
 762 probability that  $A$  is incorrect on  $C'$  is at least  $\frac{e^{-cT}}{4}$ , which contradicts the claim that  $A$  succeeds  
 763 with probability  $1 - \frac{e^{-cT}2^{-T}}{8}$ . Thus it follows that  $\sum_{j=1}^T p_j = \Omega(M)$ , as desired.  $\square$

764 Now we combine [Lemma D.9](#) and [Lemma D.10](#). This implies that the information cost can be lower  
 765 bounded by

$$I(\Pi; C_1, C_2, \dots, C_T) \geq \Omega\left(\sum_{j=1}^T p_j\right) \geq \gamma M, \quad (\text{D.8})$$

766 for a constant  $\gamma > 0$ . This completes the proof.  $\square$

## 767 E An Alternative Proof in the Large Mistake Regime

768 We give another analysis of the information cost when  $M = \Omega(T)$ , where  $M$  is the number of  
 769 mistakes of the best expert.

770 **Lemma E.1** (Single-Coordinate Information Cost Lower Bound). *Let  $c \in (0, 1)$  and  $\Pi$  be any*  
 771 *protocol with error  $\delta = 2^{-T}$  for  $\text{BIASDETECT}_1$ . Suppose that the best expert makes  $M = c'T$*   
 772 *mistakes for some constant  $c'$ . We have that the information cost of  $\Pi$  with respect to  $\zeta$  is at least*

$$I(\Pi(C_1, \dots, C_T); C_1, \dots, C_T) \geq \Omega((1-c)^2 T), \quad (\text{E.1})$$

773 where  $C_i \sim \zeta$  are i.i.d. single coordinates.

774 Applying direct sum theorem ([Lemma D.5](#)), we get the following information complexity lower  
 775 bound for  $\text{BIASDETECT}_n$ :

**Theorem E.2** ( $n$ -Coordinate Information Complexity Lower Bound). *Let  $c \in (0, 1)$  and assume*  
 *$M = c'T$  for some constant  $c'$ . Then*

$$\text{IC}_{\mathcal{D}(1), 2^{-\Theta(T)}}(\text{BIASDETECT}_n) = \Omega((1-c)^2 nT).$$

776 By an argument similar to [Theorem D.8](#), we have:

777 **Theorem E.3** (Memory lower bound for expert learning). *Let  $M = c'T$  for some constant  $c'$ . Any*  
 778 *streaming algorithm that achieves constant regret for the experts problem with probability at least*  
 779  *$1 - 2^{-\Theta(T)}$  must use at least  $\Omega(n)$  space, where the best expert makes  $M$  mistakes.*

780 For the purpose of proving [Lemma E.1](#), we need some technical lemmas.

781 **Lemma E.4** (Lemma 3.5 of [39]). *Consider any communication protocol  $\Pi$  where each player*  
 782 *receives one bit and condition on any fixed input  $b \in \{0, 1\}^T$ . Each player  $i$  can be implemented*  
 783 *such that, if the other players receive input  $b_{-i}$ , player  $i$  only observes their input with probability*  
 784  *$d_{TV}(\Pi_b, \Pi_{b \oplus e_i})$ .*

785 **Lemma E.5** (Lemma 3.6 of [39]). *Let  $c \in (0, 1)$ ,  $p \in (0, \frac{1-c}{2})$  and  $\gamma_c = \frac{1}{c \log(e/c)}$ . For a set of*  
 786 *binary random variables  $Y_1, Y_2, \dots, Y_k$  such that  $\mathbb{E}[\sum_i Y_i] = pk$ , there exists a set  $S \subset [n]$  of size*  
 787  *$ck$  such that  $\Pr(Y_j = 0, \forall j \in S) > e^{-k/\gamma_c - 1}$ .*

788 *Proof of Lemma E.1.* Let  $(C_1, C_2, \dots, C_n) \sim \zeta^n$  be a single column drawn from the NO case,  
 789 where each coordinate is i.i.d. Bernoulli with parameter  $1/2$ . Let  $M = c'T$  for some constant  $c'$ . We  
 790 consider the one-way message-passing model, where for all  $i < T$ ,  $M_i$  denotes the message sent  
 791 from player  $P_i$  to player  $P_{i+1}$ . It suffices to lower bound

$$I(\Pi; C_1, \dots, C_T) = \sum_{j=1}^T I(\Pi; C_j | C_{<j}).$$

792 by the chain rule of mutual information. We claim that for any  $j$

$$I(\Pi; C_j | C_{<j}) = I(\Pi; C_j | C_{-j}).$$

793 First, by data processing and the one-way nature of the protocol

$$I(\Pi; C_j | C_{<j}) = I(M_{\leq j}; C_j | C_{<j}).$$

794 for any  $j$ . Now we just need to show that

$$I(M_{\leq j}; C_j | C_{<j}) = I(\Pi; C_j | C_{<j}).$$

795 By chain rule of mutual information, we can write the right-hand side as

$$\begin{aligned} I(\Pi; C_j | C_{<j}) &= I(M_{\leq j}; C_j | C_{<j}) + I(M_{>j}; C_j | M_{\leq j}, C_{<j}) \\ &= I(M_{\leq j}; C_j | C_{<j}) + I(M_{>j}; C_j | M_{\leq j}, C_{>j}) \end{aligned}$$

796 Observe that  $M_{>j}$  and  $C_j$  are independent, conditional on  $M_{\leq j}$  and  $C_{>j}$ . Hence,

$$I(M_{>j}; C_j | M_{\leq j}, C_{>j}) = 0$$

797 and this proves the claim.

798 Let  $\Pi_b$  be the distribution of the protocol transcript when the input is fixed to be  $b \in \{0, 1\}^n$  and  $\oplus$   
799 denote the binary XOR. Now we can bound

$$\begin{aligned} I(\Pi; C_1, \dots, C_T) &= \sum_{j=1}^T I(\Pi; C_j | C_{<j}) \\ &= \sum_{j=1}^T I(\Pi; C_j | C_{<j}) \\ &\geq \frac{1}{8} \frac{1}{2^T} \sum_{b \in \{0,1\}^T} \sum_{j=1}^T d_{\text{TV}}^2(\Pi_{b \oplus e_j}, \Pi_b) \\ &\geq \frac{1}{8} \frac{1}{2^T} \sum_{b \in \{0,1\}^T} \sum_{j: b_j=0} d_{\text{TV}}^2(\Pi_{b \oplus e_j}, \Pi_b). \end{aligned} \quad (\text{E.2})$$

800 Conditioned on an input  $b \in \{0, 1\}^T$ , let  $k = |\{i : b_i = 0\}|$  and assume for the sake of a contradiction  
801 that

$$\sum_{i: b_i=0} d_{\text{TV}}(\Pi_{b \oplus e_i}, \Pi_b) = kp, \quad (\text{E.3})$$

802 where  $p < \frac{1-c}{2}$ . Let  $p_i = d_{\text{TV}}(\Pi_{b \oplus e_i}, \Pi_b)$  for every player  $i \in [T]$ . [Lemma E.4](#) implies that the  
803 protocol can be equivalently implemented such that if the other players receive  $b_{-i}$ , player  $i$  only  
804 looks at their input with probability  $p_i$ . If the player  $i$  does not look at their bit, then their message  
805  $M_i$  is independent of their input bit. Let  $Y_i$  denote the indicator random variable for the event that  
806 player  $i$  looks at their input in this equivalent protocol.

807 It follows from our assumption (E.3) that if the input is  $b$ , then  $\mathbb{E}[\sum_{i: b_i=0} Y_i] = \sum_i p_i = kp$ . By  
808 the definition of  $Y_i$ , if for any set  $S$ ,  $Y_i = 0$  for all  $i \in S$ , then all players in  $S$  do not look at their  
809 input bits. Let  $E_S$  denotes the event that  $Y_i = 0$  for all  $i \in S$ , for some  $S \subseteq \{i : b_i = 0\}$ . Then since  
810 the players in  $S$  do not look at their input bits,

$$d_{\text{TV}}(\Pi_{b \oplus e_S} | E_S, \Pi_b | E_S) = 0.$$

811 In particular, using this and the law of total probability, we get that

$$\begin{aligned} d_{\text{TV}}(\Pi_{b \oplus e_S}, \Pi_b) &= \Pr(E_S) \cdot d_{\text{TV}}(\Pi_{b \oplus e_S} | E_S, \Pi_b | E_S) + \Pr(\overline{E_S}) \cdot d_{\text{TV}}(\Pi_{b \oplus e_S} | \overline{E_S}, \Pi_b | \overline{E_S}) \\ &\leq \Pr(\overline{E_S}). \end{aligned} \quad (\text{E.4})$$

812 By our assumption,  $\mathbb{E}[\sum_{i: b_i=0} Y_i] = kp$  for  $p < \frac{1-c}{2}$ . Applying [Lemma E.5](#), we obtain that  
813 there exists a set  $S \subseteq \{i : b_i = 0\}$  with  $|S| = ck$  such that  $\Pr(E_S) \geq e^{-k/\gamma_c - 1}$ . For any  
814  $k < (T-2)\gamma_c < T-2$ , we have  $\Pr(E_S) > e\delta$ , and so  $\Pr(\overline{E_S}) < 1 - e\delta$ . By Eqn. (E.4),  
815  $d_{\text{TV}}(\Pi_{b \oplus e_S}, \Pi_b) < 1 - e\delta$ . Observe that  $b \oplus e_S$  differs from  $b$  by having  $|S| = ck$  more 1's; and they  
816 have same value at all other coordinates. Recall that in a typical single-coordinate YES instance, there  
817 are  $T - M$  number of 1's, which is  $T/2 - M$  more than a typical NO instance. Now suppose this

818 gap  $T/2 - M < ck$ ; then solving  $\text{BIASDETECT}_1$  is at most as hard as distinguishing  $b$  and  $b \oplus e_S$ .  
 819 Hence, if we choose  $c'$  such that  $M = c'T > T/2 - ck$ , then the protocol  $\Pi$  fails with probability  
 820 greater than  $\delta$ . This is a contradiction.

821 Thus, for any  $b$  such that  $ck = c \cdot |\{i : b_i = 0\}| > T/2 - M$ ,

$$\sum_{i:b_i=0} d_{\text{TV}}(\Pi_{b \oplus e_i}, \Pi_b) \geq \Omega\left(\frac{(1-c)T}{2}\right).$$

822 From (E.2) and Jensen's inequality,

$$I(\Pi; C_1, \dots, C_T) \geq \Omega((1-c)^2 T).$$

823 This finishes the proof. □