
Sample Complexity of Forecast Aggregation

Tao Lin

Harvard University
Cambridge, MA 02138
tlin@g.harvard.edu

Yiling Chen

Harvard University
Cambridge, MA 02138
yiling@seas.harvard.edu

Abstract

We consider a Bayesian forecast aggregation model where n experts, after observing private signals about an unknown binary event, report their posterior beliefs about the event to a principal, who then aggregates the reports into a single prediction for the event. The signals of the experts and the outcome of the event follow a joint distribution that is unknown to the principal, but the principal has access to i.i.d. “samples” from the distribution, where each sample is a tuple of the experts’ reports (not signals) and the realization of the event. Using these samples, the principal aims to find an ε -approximately optimal aggregator, where optimality is measured in terms of the expected squared distance between the aggregated prediction and the realization of the event. We show that the sample complexity of this problem is at least $\tilde{\Omega}(m^{n-2}/\varepsilon)$ for arbitrary discrete distributions, where m is the size of each expert’s signal space. This sample complexity grows exponentially in the number of experts n . But, if the experts’ signals are independent conditioned on the realization of the event, then the sample complexity is significantly reduced, to $\tilde{O}(1/\varepsilon^2)$, which does not depend on n . Our results can be generalized to non-binary events. The proof of our results uses a reduction from the distribution learning problem and reveals the fact that forecast aggregation is almost as difficult as distribution learning.

1 Introduction

Suppose you want to know whether it will rain tomorrow. A Google search on “weather” returns 40% probability of raining. The weather forecasting app on your phone shows 85%. And one of your friends, who is an expert in meteorology, predicts 65%. How do you aggregate these different predictions into a single, accurate prediction? This problem is called *forecast aggregation* or *forecast combination* [7, 16, 45]. It has innumerable applications in fields ranging from economics, statistics, operations research, to machine learning, decision theory, and of course, climate science.

A straightforward solution to forecast aggregation is to take the (unweighted) average of all experts’ forecasts. Simple as it is, unweighted average performs surprisingly well in practice, as observed by many forecast aggregation works from 1960s to 2020s (e.g., [7, 38, 35, 30, 16, 45, 34]). Naturally, when past data of expert forecasts and event outcomes (e.g., historical weather forecasts and outcomes) are available, one may hope to leverage on such data to learn more accurate aggregators. While adjusting the weights of experts in averaging according to their individual historical accuracy has led to improved accuracy for the aggregated prediction [49], interestingly, more sophisticated data-driven methods [23, 35, 49] were often outperformed by the unweighted average. The dominance of simple aggregators over more sophisticated data-driven methods is observed so often in empirical applications that it is termed “the forecast combination puzzle” [44, p.428].

There are many potential explanations for the forecast combination puzzle [43, 23, 15]. In some scenarios, the past events are different from the future events in fundamental ways (e.g., geopolitical

forecasting) and hence past data may not be informative in predicting the future. Another widely accepted conjecture is that the amount of past data is not large enough for a data-intensive method to perform well. Indeed, the sample sizes in many empirical forecast aggregation works are small under today’s “big-data” standard (24 in [27], 69 in [42], 87 in [6], 100 in [30]). However, there are aggregation settings where future events are arguably similar to past events and we do have abundant data — for instance, the forecasting of weather, stock prices [21], and some forecasting competitions on Kaggle¹. Such settings are well-suited for data-driven methods. It is hence tempting to ask *how many data* are needed for data-driven aggregators to achieve high accuracy in these settings.

In this paper, we initiate the study of *sample complexity of forecast aggregation*, building upon a standard Bayesian forecasting model [37, 48, 24]. In this model, there are n experts who observe private signals about an unknown binary event and make posterior predictions about the event. The experts’ signals and the event are jointly distributed according to some underlying unknown distribution (information structure) P , which determines the optimal way to aggregate the experts’ predictions. With sample access to the unknown distribution P , we ask the following question:

How many samples do we need to approximately learn the optimal aggregator?

Our model favors the use of data-driven aggregation methods because historical tasks are i.i.d. as future tasks. We show that however, even in this benign model, optimal aggregation in general needs *exponentially many* samples. One may thus expect that data-driven methods can hardly perform well in more realistic and non-i.i.d. scenarios. Nevertheless, for some special, yet interesting, families of information structures, the sample complexity of forecast aggregation is significantly lower.

Main results (1) If P can be an arbitrary discrete distribution, then at least $\tilde{\Omega}(m^{n-2}/\varepsilon)$ samples are needed, and $\tilde{O}(m^n/\varepsilon^2)$ samples are sufficient, to learn an ε -optimal aggregator with high probability, where m is the number of signals an expert can possibly observe.² (2) If the experts’ signals are conditionally independent, then the sample complexity is exponentially reduced and surprisingly does not depend on the number of experts or the number of signals: $\tilde{O}(1/\varepsilon^2)$ samples are sufficient, and $\tilde{\Omega}(1/\varepsilon)$ samples are necessary, to learn an ε -optimal aggregator with high probability.

Main techniques The main technical part of our paper is to prove the $\tilde{\Omega}(m^{n-2}/\varepsilon)$ lower bound on the sample complexity of forecast aggregation for the general case, via a reduction from the *distribution learning* problem. It is known that learning a discrete distribution with support size $|S|$ in total variation distance ε requires $\tilde{\Omega}(|S|/\varepsilon^2)$ samples. By reducing distribution learning to forecast aggregation, we obtain the lower bound on the sample complexity of the latter problem. This reduction is highly nontrivial. To do this reduction we define a new distribution learning problem that is different from the one in the literature, which is of independent interest. This reduction also reveals an interesting fact: learning to aggregate optimally on some distribution is almost as difficult as learning the distribution itself. This is a little surprising because one might initially think that aggregation should be easier than distribution learning – we show that this is not the case.

Structure of the paper We discuss related works in Section 1.1. Section 2 introduces our model. Section 3 includes some preliminaries, including the distribution learning problem, which we will use in the proof. Section 4 studies the sample complexity for general distributions. Section 5 focuses on the conditional independence case. Section 6 summarizes how our results can be generalized to non-binary (multi-outcome) events. Section 7 concludes and discusses future directions.

1.1 Related Works

Data-driven aggregation Data-driven approaches to forecast aggregation date back to perhaps the first paper on forecast aggregation [7] and have been a standard practice in the literature (see, e.g., surveys [16, 45] and more recent works [51, 42, 4, 47, 22, 40]). Many of these works focus on specific weighted average aggregators like *linear pooling* (weighted arithmetic mean) [51] and *logarithmic pooling* (normalized weighted geometric mean) [42, 40], and the goal is to estimate the optimal weights from data. However, those weighted average aggregators are not necessarily the

¹<https://www.kaggle.com/c/m5-forecasting-accuracy>

²The $\tilde{O}(\cdot)$ and $\tilde{\Omega}(\cdot)$ notations omit logarithmic factors.

optimal (*Bayesian*) aggregator unless the underlying information structure satisfies some strict conditions (e.g., experts’ forecasts are equal to the true probability of the event plus normally distributed noises [48]). Our work aims to understand the sample complexity of Bayesian aggregation, namely how many samples are needed to approximate the Bayesian aggregator.

Our work is closely related to a recent work [4] on Bayesian aggregation via *online learning*. For the case of conditionally independent experts, [4] shows that the Bayesian aggregator can be approximated with $\varepsilon = \tilde{O}(\frac{n}{\sqrt{T}})$ regret. By an online-to-batch conversion, this implies a $T = \tilde{O}(\frac{n^2}{\varepsilon^2})$ sample complexity upper bound for the batch learning setting. Our paper studies the batch learning setting. For the conditional independence case, we obtain an improved bound of $\tilde{O}(\frac{1}{\varepsilon^2})$.

Robust forecast aggregation Recent works on “robust forecast aggregation” [2, 39, 18, 3] also study information aggregation problems where the principal does not know the underlying information structure. They take a worst-case approach, assuming that the information structure is chosen adversarially. This often leads to negative results: e.g., a bad approximation ratio [2, 39] or a degenerate maximin aggregator that solely relies on a random expert’s opinion [18, 3]. In contrast, we assume sample access to the unknown information structure. Our sample complexity approach is orthogonal and complementary to the robust forecast aggregation approach.

Sample complexity of mechanism design Our work may remind the reader of a long line of research on the sample complexity of revenue maximization in mechanism design (e.g., [17, 20, 36, 19, 5, 11, 29, 10, 25, 50, 28]). In particular, [28] gives a general framework to bound the sample complexity for mechanism design problems that satisfy a “strong monotonicity” property, but this property is not satisfied by our forecast aggregation problem. A key observation from this line of works is that the number of samples needed to learn an ε -optimal auction for $n \geq 1$ independent bidders is increasing in n , because when there are more bidders, although we can obtain higher revenue, the optimal auction benchmark is also improved. We see a similar phenomenon that the sample complexity of forecast aggregation increases with the number of experts in the general case, but not in the case where experts are conditionally independent.

2 Model

2.1 Forecast Aggregation

There are $n \geq 2$ experts and one principal. The principal wants to predict the probability that a binary event $\omega \in \Omega = \{0, 1\}$ happens ($\omega = 1$), based on information provided by the experts. For example, ω may represent whether it will rain tomorrow. We present binary events to simplify notations. All our results can be generalized to multi-outcome events with $|\Omega| > 2$ (see Section 6). We also refer to ω as “the state of the world”. Each expert $i = 1, \dots, n$ observes some private signal $s_i \in \mathcal{S}_i$ that is correlated with ω , where \mathcal{S}_i denotes the space of all possible signals of expert i . We assume for now that \mathcal{S}_i is finite, with size $|\mathcal{S}_i| = m$. We relax this assumption in Section 5 where we consider conditionally independent signals. Let $\mathcal{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ be the joint signal space of all experts; $|\mathcal{S}| = m^n$. Let P be a distribution over $\mathcal{S} \times \Omega$, namely, a joint distribution of signals $\mathbf{s} = (s_1, \dots, s_n)$ and event ω . Since the space $\mathcal{S} \times \Omega$ is discrete, we can use $P(\cdot)$ to denote the probability: $P(\mathbf{s}, \omega) = \Pr_P[\mathbf{s}, \omega]$. Signals of different experts can be correlated conditioned on ω . We assume that each expert i knows the marginal joint distribution of their own signal s_i and ω , $P(s_i, \omega)$. Neither any expert nor the principal knows the entire distribution P . Each expert i reports to the principal a forecast (or prediction) r_i for the event ω , which is equal to the conditional probability of $\omega = 1$ given their signal s_i :³

$$r_i = P(\omega = 1 \mid s_i) = \frac{P(\omega=1)P(s_i|\omega=1)}{P(\omega=1)P(s_i|\omega=1)+P(\omega=0)P(s_i|\omega=0)}. \quad (1)$$

We note that r_i depends on s_i and P , but not on ω or other experts’ signals \mathbf{s}_{-i} . Let $\mathbf{r} = (r_1, \dots, r_n) \in [0, 1]^n$ denote the reports (joint report) of all experts. We sometimes use \mathbf{r}_{-i} to

³One may wonder whether the experts are willing to report $r_i = P(\omega = 1 \mid s_i)$ truthfully. This can be guaranteed by a *proper scoring rule*. For example, we can reward each expert the Brier score $C - |r_i - \omega|^2$ after seeing the realization of ω [9]. Each expert maximizes its expected reward by reporting its belief truthfully.

denote the reports of all experts except i . The principal aggregates the experts' reports \mathbf{r} into a single forecast $f(\mathbf{r})$ using some *aggregation function*, or *aggregator*, $f : [0, 1]^n \rightarrow [0, 1]$. We measure the performance of an aggregator by the mean squared loss:

$$L_P(f) = \mathbb{E}_P[|f(\mathbf{r}) - \omega|^2]. \quad (2)$$

The notation $\mathbb{E}_P[\cdot]$ makes it explicit that the expectation is over the random draw of $(\mathbf{s}, \omega) \sim P$ followed by letting $r_i = P(\omega = 1 | s_i)$. We omit P and write $\mathbb{E}[\cdot]$ when it is clear from the context.

Let f^* be the optimal aggregator with respect to P , which minimizes $L_P(f)$:

$$f^* = \operatorname{argmin}_{f: [0,1]^n \rightarrow [0,1]} L_P(f) = \operatorname{argmin}_{f: [0,1]^n \rightarrow [0,1]} \mathbb{E}_P[|f(\mathbf{r}) - \omega|^2]. \quad (3)$$

We have the following characterization of f^* and $L_P(f)$: f^* is equal to the Bayesian aggregator, which computes the posterior probability of $\omega = 1$ given all the reports $\mathbf{r} = (r_1, \dots, r_n)$. And the difference between the loss of f and the loss of f^* is equal to their expected squared difference.

Lemma 2.1. *The optimal aggregator f^* and any aggregator f satisfy:*

- $f^*(\mathbf{r}) = P(\omega = 1 | \mathbf{r})$, for almost every \mathbf{r} .
- $L_P(f) - L_P(f^*) = \mathbb{E}_P[|f(\mathbf{r}) - f^*(\mathbf{r})|^2]$.

An aggregator f is ε -optimal (with respect to P) if $L_P(f) \leq L_P(f^*) + \varepsilon$. By Lemma 2.1, this is equivalent to $\mathbb{E}_P[|f(\mathbf{r}) - f^*(\mathbf{r})|^2] \leq \varepsilon$. For an ε -optimal f , we also say it ε -approximates f^* .

Definition 2.2. *An aggregator f is ε -optimal (with respect to P) if $\mathbb{E}_P[|f(\mathbf{r}) - f^*(\mathbf{r})|^2] \leq \varepsilon$.*

Discussion of the benchmark f^* Our benchmark, the Bayesian aggregator f^* , is common in the forecast aggregation literature (e.g., [27, 26, 45]). It is stronger than the typical ‘‘best expert’’ benchmark in no-regret learning (e.g., [13, 22]), but weaker than the ‘‘omniscient’’ aggregator that has access to the experts’ signals: $f_{\text{omni}}(\mathbf{s}) = P(\omega = 1 | \mathbf{s})$. If there is a one-to-one mapping between signals \mathbf{s} and reports \mathbf{r} , then f_{omni} and f^* are the same. Otherwise, f_{omni} could be much stronger than f^* and an ε -approximation to f_{omni} using experts’ reports only is not always possible.⁴ In contrast, an ε -approximation to f^* is always achievable (in fact, achieved by f^* itself). The difference between f^* and f_{omni} is known as the difference between ‘‘aggregating forecasts’’ and ‘‘aggregating information sets’’ in the literature [27, p.198-199], [26, p.168-169], [45, p.143].

2.2 Sample Complexity of Forecast Aggregation

The principal has access to T i.i.d. samples of forecasts and event realizations drawn from the underlying unknown distribution P :

$$S_T = \{(\mathbf{r}^{(1)}, \omega^{(1)}), \dots, (\mathbf{r}^{(T)}, \omega^{(T)})\}, \quad (\mathbf{s}^{(t)}, \omega^{(t)}) \sim P, \quad r_i^{(t)} = P(\omega = 1 | s_i^{(t)}). \quad (4)$$

Here, we implicitly regard P as a distribution over (\mathbf{r}, ω) instead of (\mathbf{s}, ω) . The principal uses samples S_T to learn an aggregator $\hat{f} = \hat{f}_{S_T}$, in order to approximate f^* . Our main question is:

How many samples are necessary and sufficient for finding an ε -optimal aggregator \hat{f} (with probability at least $1 - \delta$ over the random draw of samples)?

The answer to the above question depends on the family of distributions we are interested in. Let \mathcal{P} denote a family of distributions over $\mathcal{S} \times \Omega$. It could be the set of all distributions over $\mathcal{S} \times \Omega$, or in Section 5 we will only consider distributions where the signals are independent conditioned on ω . We define the sample complexity of forecast aggregation, with respect to \mathcal{P} , formally:

Definition 2.3. *The sample complexity of forecast aggregation (with respect to \mathcal{P}) is the minimum function $T_{\mathcal{P}}(\cdot, \cdot)$ of $\varepsilon, \delta \in (0, 1)$, such that: if $T \geq T_{\mathcal{P}}(\varepsilon, \delta)$, then for any distribution $P \in \mathcal{P}$, with probability at least $1 - \delta$ over the random draw of T samples S_T from P (and over the randomness of the learning procedure if it is randomized), we can obtain an aggregator $\hat{f} = \hat{f}_{S_T}$ satisfying $\mathbb{E}_P[|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|^2] \leq \varepsilon$.*

⁴This has been noted by [2, 4, 39]. They give an XOR example where $\omega = s_1 \oplus s_n$, s_1 and s_2 are i.i.d. Uniform $\{0, 1\}$ distributed. The experts always report $r_i = 0.5$, $f_{\text{omni}}(s_1, s_2) = \omega$, $f^*(r_1, r_2) = 0.5$, so $L_P(f_{\text{omni}}) = 0$ but $L_P(f^*) = 0.25 > 0$. No aggregator that uses experts’ reports only can do better than f^* .

The principal is assumed to know the family of distributions \mathcal{P} but not the specific distribution $P \in \mathcal{P}$. There should be at least two different distributions in \mathcal{P} . Otherwise, the principal knows what the distribution is and there is no need for learning.

3 Preliminaries

In this section, we briefly introduce some notions that will be used in our analysis of the sample complexity of forecast aggregation, including some definitions of distances between distributions, the distribution learning problem, and the distinguishing distributions problem.

Distances between distributions We recall two distance metrics for discrete distributions: the *total variation distance* and the (*squared*) *Hellinger distance*.

Definition 3.1. Let D_1, D_2 be two distributions on a discrete space \mathcal{X} .

- The total variation distance between D_1 and D_2 is $d_{\text{TV}}(D_1, D_2) = \frac{1}{2} \sum_{x \in \mathcal{X}} |D_1(x) - D_2(x)|$.
- The squared Hellinger distance between D_1 and D_2 is $d_{\text{H}}^2(D_1, D_2) = \frac{1}{2} \sum_{x \in \mathcal{X}} (\sqrt{D_1(x)} - \sqrt{D_2(x)})^2 = 1 - \sum_{x \in \mathcal{X}} \sqrt{D_1(x)D_2(x)}$.

The total variation distance has the following well-known property that upper bounds the difference between the expected values of a function on two distributions:

Fact 3.2. For any function $h : \mathcal{X} \rightarrow [0, 1]$, $|\mathbb{E}_{x \sim D_1} h(x) - \mathbb{E}_{x \sim D_2} h(x)| \leq d_{\text{TV}}(D_1, D_2)$.

In Appendix A we give some properties of the Hellinger distance that will be used in the proofs.

Distribution learning in total variation distance Our analysis of the sample complexity of the forecast aggregation problem will leverage on the sample complexity of another learning problem: *learning discrete distributions in total variation distance*. We review this problem below.

Let \mathcal{D} be a family of distributions over \mathcal{X} . The *sample complexity of learning distributions in \mathcal{D} within total variation distance ε* is the minimum function $T_{\mathcal{D}}^{\text{TV}}(\varepsilon, \delta)$, such that: if $T \geq T_{\mathcal{D}}^{\text{TV}}(\varepsilon, \delta)$, then for any distribution $D \in \mathcal{D}$, with probability at least $1 - \delta$ over the random draw of T samples from D , we can obtain (from the T samples) a distribution \hat{D} such that $d_{\text{TV}}(\hat{D}, D) \leq \varepsilon$.

Proposition 3.3 (e.g., [12, 33]). Let \mathcal{D}_{all} be the set of all distributions over \mathcal{X} . Then, $T_{\mathcal{D}_{\text{all}}}^{\text{TV}}(\varepsilon, \delta) = \Theta\left(\frac{|\mathcal{X}| + \log(1/\delta)}{\varepsilon^2}\right)$. In particular, the upper bound can be achieved by using the empirical estimate \hat{D}_{emp} (which is the uniform distribution over the T samples). The lower bound holds regardless of what learning algorithm is used.

Distinguishing distributions Another learning problem that we will leverage on is the problem of *distinguishing (two) distributions*: given samples from a distribution randomly chosen from $\{D_1, D_2\}$, we are to guess whether the samples are from D_1 or D_2 . The sample complexity of distinguishing distributions is characterized by the squared Hellinger distance. It is known that at least $T = \Omega\left(\frac{1}{d_{\text{H}}^2(D_1, D_2)} \log \frac{1}{\delta}\right)$ samples are needed to distinguish two distributions with probability at least $1 - \delta$. See Appendix A for a formal statement of this result.

4 Sample Complexity for General Distributions

In this section we characterize the sample complexity of forecast aggregation for general distributions P . We give an exponential (in the number of experts, n) upper bound and an exponential lower bound on the sample complexity, as follows:

Theorem 4.1. Let \mathcal{P}_{all} be the set of all distributions over $\mathcal{S} \times \Omega$, with $|\mathcal{S}| = m^n$. Suppose $n \geq 2$. The sample complexity of forecast aggregation with respect to \mathcal{P}_{all} is

$$O\left(\frac{m^n + \log(1/\delta)}{\varepsilon^2}\right) \geq T_{\mathcal{P}_{\text{all}}}(\varepsilon, \delta) \geq \Omega\left(\frac{m^{n-2} + \log(1/\delta)}{\varepsilon}\right). \quad (5)$$

This theorem is for $n \geq 2$. When there is only one expert ($n = 1$), there is no need to learn to aggregate because the optimal ‘‘aggregator’’ f^* simply outputs the forecast given by the only expert: $f^*(r_1) = P(\omega = 1 | r_1) = P(\omega = 1 | s_1) = r_1$. The sample complexity is 0 in this case.

There is a gap in the dependency on ε in the upper bound and the lower bound in Theorem 4.1. We conjecture that the tight dependency on ε should be $\frac{1}{\varepsilon}$ (so the lower bound is tight). See Section 7 for a detailed discussion of this conjecture, where we show that the dependency on ε in the upper bound can be improved to $\frac{1}{\varepsilon}$ for a large family of distributions.

4.1 Proof of the Upper Bound

In this subsection we prove the $O\left(\frac{m^n + \log(1/\delta)}{\varepsilon^2}\right)$ upper bound in Theorem 4.1. This is a direct corollary of the distribution learning result introduced in Section 3.

We regard P as a distribution over \mathbf{r} and ω instead of over \mathbf{s} and ω . Then P is a discrete distribution with support size at most $2m^n$ because each possible report $r_i \in [0, 1]$ corresponds to some discrete signal s_i in \mathcal{S}_i . Let \hat{P}_{emp} be the empirical distribution of reports and event realizations: $\hat{P}_{\text{emp}} = \text{Uniform}\{(\mathbf{r}^{(1)}, \omega^{(1)}), \dots, (\mathbf{r}^{(T)}, \omega^{(T)})\}$. By Proposition 3.3, with probability at least $1 - \delta$ over the random draw of $T = O\left(\frac{2m^n + \log(1/\delta)}{\varepsilon^2}\right)$ samples, we have $d_{\text{TV}}(\hat{P}_{\text{emp}}, P) \leq \varepsilon$. According to Fact 3.2, and by the definition of $L_P(f)$, we have: for any aggregator $f : [0, 1]^n \rightarrow [0, 1]$,

$$|L_{\hat{P}_{\text{emp}}}(f) - L_P(f)| = \left| \mathbb{E}_{\hat{P}_{\text{emp}}} [|f(\mathbf{r}) - \omega|^2] - \mathbb{E}_P [|f(\mathbf{r}) - \omega|^2] \right| \leq d_{\text{TV}}(\hat{P}_{\text{emp}}, P) \leq \varepsilon.$$

Therefore, if we pick the empirically optimal aggregator $\hat{f}_{\text{emp}} = \text{argmin}_f L_{\hat{P}_{\text{emp}}}(f)$, we get

$$L_P(\hat{f}_{\text{emp}}) \leq L_{\hat{P}_{\text{emp}}}(\hat{f}_{\text{emp}}) + \varepsilon \leq L_{\hat{P}_{\text{emp}}}(f^*) + \varepsilon \leq L_P(f^*) + 2\varepsilon,$$

which means that \hat{f}_{emp} is a 2ε -optimal aggregator for P .

4.2 Proof of the Lower Bound

In this subsection we prove the $\Omega\left(\frac{m^{n-2} + \log(1/\delta)}{\varepsilon}\right)$ lower bound in Theorem 4.1. The main idea is a reduction from the distribution learning problem (defined in Section 3) for a specific family \mathcal{D} of distributions over the joint signal space $\mathcal{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_n$. We construct a corresponding family of distributions $\mathcal{P} = \{P_D : D \in \mathcal{D}\}$ for the forecast aggregation problem, such that, if we can obtain an ε -optimal aggregator \hat{f} for P_D , then we can convert \hat{f} into a distribution \hat{D} such that $d_{\text{TV}}(\hat{D}, D) \leq O(\sqrt{\varepsilon})$. We then prove that learning \mathcal{D} within total variation distance $\varepsilon_{\text{TV}} = O(\sqrt{\varepsilon})$ requires $\Omega\left(\frac{m^{n-2} + \log(1/\delta)}{\varepsilon_{\text{TV}}}\right) = \Omega\left(\frac{m^{n-2} + \log(1/\delta)}{\varepsilon}\right)$ samples. This gives the sample complexity lower bound for the forecast aggregation problem for \mathcal{P} (and hence \mathcal{P}_{all}).

We will need a family of distributions \mathcal{D} that satisfies the following three properties:

Definition 4.2. *We say a family of distributions \mathcal{D} satisfies*

1. *B-uniformly bounded, if: $D(\mathbf{s}) \leq \frac{B}{|\mathcal{S}|} = \frac{B}{m^n}, \forall \mathbf{s} \in \mathcal{S}, \forall D \in \mathcal{D}$, where $B \geq 1$ is a constant.*
2. *same marginal across distributions, if: for any $D, D' \in \mathcal{D}$, any i , any $s_i \in \mathcal{S}_i$, $D(s_i) = D'(s_i)$.*
3. *distinct marginals across signals, if: for any $D \in \mathcal{D}$, any i , any $s_i \neq s'_i \in \mathcal{S}_i$, $D(s_i) \neq D(s'_i)$.*

How do we construct the family \mathcal{P} ? For each distribution $D \in \mathcal{D}$, we construct distribution P_D as follows: the marginal distribution of ω is $\text{Uniform}\{0, 1\}$, i.e., $P_D(\omega = 0) = P_D(\omega = 1) = \frac{1}{2}$; conditioning on $\omega = 0$, the joint signal \mathbf{s} is uniformly distributed: $P_D(\mathbf{s} \mid \omega = 0) = \frac{1}{|\mathcal{S}|} = \frac{1}{m^n}$, $\forall \mathbf{s} \in \mathcal{S}$; conditioning on $\omega = 1$, the joint signal is distributed according to D : $P_D(\mathbf{s} \mid \omega = 1) = D(\mathbf{s}), \forall \mathbf{s} \in \mathcal{S}$. The family \mathcal{P} is $\{P_D : D \in \mathcal{D}\}$.

We show that if we can obtain ε -optimal aggregators for distributions in \mathcal{P} , then we can learn the distributions in \mathcal{D} within total variation distance $(1+B)^2\sqrt{\varepsilon}$, and thus the sample complexity of the former is lower bounded by the sample complexity of the latter:

Lemma 4.3. *Let \mathcal{D} be a family of distributions that satisfies the three properties in Definition 4.2. Let $\mathcal{P} = \{P_D : D \in \mathcal{D}\}$ be defined above. Then, $T_{\mathcal{P}}(\varepsilon, \delta) \geq T_{\mathcal{D}}^{\text{TV}}((1+B)^2\sqrt{\varepsilon}, \delta)$.*

Proof sketch of Lemma 4.3. The full proof is in Appendix E.1. We give a sketch here. According to the definition of P_D , by Bayes' rule, we have

$$P_D(\omega = 1 \mid \mathbf{s}) = \frac{\frac{1}{2}P_D(\mathbf{s}|\omega=1)}{\frac{1}{2}P_D(\mathbf{s}|\omega=0) + \frac{1}{2}P_D(\mathbf{s}|\omega=1)} = \frac{D(\mathbf{s})}{\frac{1}{m^n} + D(\mathbf{s})}. \quad (6)$$

The ‘‘distinct marginals across signals’’ property in Definition 4.2 ensures that there is a one-to-one mapping between signal s_i and report $r_i = P_D(\omega = 1 \mid s_i)$, and hence a one-to-one mapping between joint signal \mathbf{s} and joint report \mathbf{r} . So, the Bayesian aggregator f^* satisfies $f^*(\mathbf{r}) = P_D(\omega = 1 \mid \mathbf{r}) = P_D(\omega = 1 \mid \mathbf{s}) = \frac{D(\mathbf{s})}{1/m^n + D(\mathbf{s})}$. This gives

$$D(\mathbf{s}) = \frac{1}{m^n} \frac{f^*(\mathbf{r})}{1 - f^*(\mathbf{r})}. \quad (7)$$

Suppose we have obtained an ε -optimal aggregator \hat{f} for P_D , $\mathbb{E}_{P_D} [|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|^2] \leq \varepsilon$, then we convert \hat{f} into \hat{D} by letting $\hat{D}(\mathbf{s}) = \frac{1}{m^n} \frac{\hat{f}(\mathbf{r})}{1 - \hat{f}(\mathbf{r})}$. The total variation distance between \hat{D} and D is:

$$d_{\text{TV}}(\hat{D}, D) = \frac{1}{2} \sum_{\mathbf{s} \in \mathcal{S}} |\hat{D}(\mathbf{s}) - D(\mathbf{s})| \stackrel{(7)}{=} \frac{1}{2} \sum_{\mathbf{s} \in \mathcal{S}} \frac{1}{m^n} \left| \frac{\hat{f}(\mathbf{r})}{1 - \hat{f}(\mathbf{r})} - \frac{f^*(\mathbf{r})}{1 - f^*(\mathbf{r})} \right|. \quad (8)$$

The ‘‘ B -uniformly bounded’’ property in Definition 4.2 ensures $D(\mathbf{s}) = O(\frac{1}{m^n})$, which has two consequences: (1) $P_D(\mathbf{s}) = O(\frac{1}{m^n})$; (2) $f^*(\mathbf{r}) = O(1)$, which implies $\left| \frac{\hat{f}(\mathbf{r})}{1 - \hat{f}(\mathbf{r})} - \frac{f^*(\mathbf{r})}{1 - f^*(\mathbf{r})} \right| = O(|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|)$ due to Lipschitz property of the function $\frac{x}{1-x}$ for $x = O(1)$. These two consequences imply

$$\begin{aligned} d_{\text{TV}}(\hat{D}, D) &= O\left(\sum_{\mathbf{s} \in \mathcal{S}} \frac{1}{m^n} |\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|\right) = O\left(\sum_{\mathbf{s} \in \mathcal{S}} P_D(\mathbf{s}) |\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|\right) \\ &= O\left(\mathbb{E}[|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|]\right) \stackrel{\text{Jensen's inequality}}{\leq} O\left(\sqrt{\mathbb{E}[|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|^2]}\right) = O(\sqrt{\varepsilon}). \end{aligned}$$

So, we obtain $d_{\text{TV}}(\hat{D}, D) \leq O(\sqrt{\varepsilon}) = (1+B)^2 \sqrt{\varepsilon}$.

There is a subtlety, however: In the distribution learning problem we are given samples from D , which are *signals* $\{\mathbf{s}^{(t)}\}_{t=1}^T$. We need to convert them into the corresponding *reports* $\{\mathbf{r}^{(t)}\}_{t=1}^T$ as the samples for the forecast aggregation problem. To do this we need to know P_D or D , which we *do not know*. So, we make use of the ‘‘same marginal across distributions’’ property in Definition 4.2 here: because all the distributions $D \in \mathcal{D}$ have the same marginal probability $D(s_i) = D'(s_i)$ (but possibly different joint probabilities $D(\mathbf{s}) \neq D'(\mathbf{s})$), we are able to compute the report

$$r_i^{(t)} = P_D(\omega = 1 \mid s_i^{(t)}) = \frac{\frac{1}{2}P_D(s_i^{(t)}|\omega=1)}{\frac{1}{2}P_D(s_i^{(t)}|\omega=0) + \frac{1}{2}P_D(s_i^{(t)}|\omega=1)} = \frac{D(s_i^{(t)})}{\frac{1}{m} + D(s_i^{(t)})}$$

separately for each expert i without knowing D , since we know what the family \mathcal{D} is. This allows us to reduce the distribution learning problem for \mathcal{D} to the forecast aggregation problem for \mathcal{P} . \square

Then, we find a family of distributions \mathcal{D} that satisfies the three properties in Definition 4.2 and requires many samples to learn.

Proposition 4.4. *There exists a family of distributions \mathcal{D} that satisfies the three properties in Definition 4.2 (with $B = e + \frac{1}{2}$) and requires $T_{\mathcal{D}}^{\text{TV}}(\varepsilon_{\text{TV}}, \delta) = \Omega\left(\frac{m^{n-2} + \log(1/\delta)}{\varepsilon_{\text{TV}}^2}\right)$ samples to learn.*

The above sample complexity is smaller than the $\Omega\left(\frac{|\mathcal{S}| + \log(1/\delta)}{\varepsilon_{\text{TV}}^2}\right)$ lower bound in Proposition 3.3 because we are restricting to a smaller set of distributions than the set of all distributions over \mathcal{S} . The proof of Proposition 4.4 is analogous to a textbook proof of Proposition 3.3, which uses reductions from the distinguishing distributions problem. See details in Appendix E.2.

Finishing the proof of Theorem 4.1: By Lemma 4.3 and Proposition 4.4, plugging in $\varepsilon_{\text{TV}} = (1+B)^2 \sqrt{\varepsilon}$ with $B = e + \frac{1}{2}$, we obtain the lower bound on the sample complexity of forecast aggregation for \mathcal{P} (and hence for \mathcal{P}_{all}):

$$T_{\mathcal{P}}(\varepsilon, \delta) \geq T_{\mathcal{D}}^{\text{TV}}((1+B)^2 \sqrt{\varepsilon}, \delta) = \Omega\left(\frac{m^{n-2} + \log(1/\delta)}{((1+B)^2 \sqrt{\varepsilon})^2}\right) = \Omega\left(\frac{m^{n-2} + \log(1/\delta)}{\varepsilon}\right).$$

5 Sample Complexity for Conditionally Independent Distributions

Section 4 proved that learning ε -optimal aggregators for all discrete distributions needs exponentially many samples. As shown in the proof, this large sample complexity is because the experts' signals can be arbitrarily correlated conditioned on the event ω . Accurate estimation of such correlation requires many samples. So, in this section we restrict attentions to the case where the experts' signals are conditionally independent. It turns out that an ε -optimal aggregator can be learned using only $O(\frac{1}{\varepsilon^2} \log \frac{1}{\varepsilon\delta})$ samples in this case, which does not depend on n . The assumption of discrete signal space can be relaxed here. We also investigate two special and interesting families of conditionally independent distributions that admit an even smaller sample complexity of $O(\frac{1}{\varepsilon} \log \frac{1}{\delta})$.

5.1 General Conditionally Independent Distributions

Let P be a conditionally independent distributions over $\mathcal{S} \times \Omega$, namely, $P(\mathbf{s} | \omega) = \prod_{i=1}^n P(s_i | \omega)$ for all $\mathbf{s} \in \mathcal{S}$, for $\omega \in \{0, 1\}$. Here, \mathcal{S}_i can be a continuous space, in which case, $P(\cdot | \omega)$ represents a density function. We introduce some additional notations. Let $p = P(\omega = 1)$ be the prior probability of $\omega = 1$. For technical convenience we assume $p \in (0, 1)$. Define

$$\rho = \frac{p}{1-p} = \frac{P(\omega=1)}{P(\omega=0)} \in (0, +\infty). \quad (9)$$

We will be working with ratios like “ $\frac{r_i}{1-r_i}$ ” and “ $\frac{p}{1-p}$ ” a lot in this section. We will use the following characterization of the optimal aggregator f^* for conditionally independent distributions:

Lemma 5.1 ([8]). *For conditionally independent distribution P , given signals $\mathbf{s} = (s_i)_{i=1}^n$, with corresponding reports $\mathbf{r} = (r_i)_{i=1}^n$ where $r_i = P(\omega = 1 | s_i)$, the posterior probability of $\omega = 1$ is:*

$$f^*(\mathbf{r}) = P(\omega = 1 | \mathbf{r}) = P(\omega = 1 | \mathbf{s}) = \frac{1}{1 + \rho^{n-1} \prod_{i=1}^n \frac{1-r_i}{r_i}}. \quad (10)$$

(Define $f^*(\mathbf{r}) = 0$ if $\rho^{n-1} \prod_{i=1}^n \frac{1-r_i}{r_i} = +\infty$.)

Lemma 5.1 implies that one way to learn f^* is to simply learn the value of ρ . If we can learn ρ with accuracy $\frac{\sqrt{\varepsilon}}{n}$, then we can learn ρ^{n-1} with accuracy $\sqrt{\varepsilon}$ and obtain an \hat{f} that is ε -close to f^* for every possible input $\mathbf{r} \in [0, 1]^n$. However, by standard concentration inequalities, learning ρ with accuracy $\frac{\sqrt{\varepsilon}}{n}$ requires $\tilde{O}(\frac{n^2}{\rho\varepsilon})$ samples, which is larger than the $\tilde{O}(\frac{1}{\varepsilon^2})$ bound we will prove. The key is that we do not actually need $\hat{f}(\mathbf{r})$ to be close to $f^*(\mathbf{r})$ for every $\mathbf{r} \in [0, 1]^n$; we only need the expectation $\mathbb{E}[|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|^2] \leq \varepsilon$. This allows us to prove a smaller sample complexity bound, using a pseudo-dimension argument.

The main result of this section is that the sample complexity of forecast aggregation with respect to all conditionally independent distributions is between $\Omega(\frac{1}{\varepsilon} \log \frac{1}{\delta})$ and $O(\frac{1}{\varepsilon^2} \log \frac{1}{\varepsilon\delta})$:

Theorem 5.2. *Let \mathcal{P}_{ind} be the set of all conditionally independent distributions over $\mathcal{S} \times \Omega$. Suppose $n \geq 2$. The sample complexity of forecast aggregation with respect to \mathcal{P}_{ind} is*

$$O(\frac{1}{\varepsilon^2} \log \frac{1}{\varepsilon\delta}) \geq T_{\mathcal{P}_{\text{ind}}}(\varepsilon, \delta) \geq \Omega(\frac{1}{\varepsilon} \log \frac{1}{\delta}). \quad (11)$$

We provide the main ideas of the proof of Theorem 5.2 here. The upper bound $O(\frac{1}{\varepsilon^2} \log \frac{1}{\varepsilon\delta})$ is a corollary of our theorem for multi-outcome events (Theorem C.1), so we only give a sketch here. We note that, according to Lemma 5.1, the optimal aggregator has the form $f^*(\mathbf{r}) = \frac{1}{1 + \rho^{n-1} \prod_{i=1}^n \frac{1-r_i}{r_i}}$.

We consider the class of aggregators $\mathcal{F} = \{f^\theta : f^\theta(\mathbf{r}) = \frac{1}{1 + \theta^{n-1} \prod_{i=1}^n \frac{1-r_i}{r_i}}\}$ parameterized by $\theta \in (0, +\infty)$. The class of loss functions $\mathcal{G} = \{g^\theta : g^\theta(\mathbf{r}, \omega) = |f^\theta(\mathbf{r}) - \omega|^2\}$ associated with \mathcal{F} has pseudo-dimension $\text{Pdim}(\mathcal{G}) = O(1)$. By the known result (e.g., [1]) that the pseudo-dimension gives a sample complexity upper bound on the uniform convergence of a class of functions, we conclude that the empirically optimal aggregator in \mathcal{F} must be $O(\varepsilon)$ -optimal on the true distribution (with probability at least $1 - \delta$), given $O(\frac{1}{\varepsilon^2} (\text{Pdim}(\mathcal{G}) \log \frac{1}{\varepsilon} + \log \frac{1}{\delta})) = O(\frac{1}{\varepsilon^2} \log \frac{1}{\varepsilon\delta})$ samples.

We prove the $\Omega(\frac{1}{\varepsilon} \log \frac{1}{\delta})$ lower bound by a reduction from the distinguishing distributions problem (introduced in Section 3). We construct two conditionally independent distributions P^1, P^2 over

the space $\Omega \times \mathcal{S}$ that differ by $d_{\text{H}}^2(P^1, P^2) = O(\varepsilon)$ in squared Hellinger distance. Specifically, P^1 has prior $P^1(\omega = 1) = 0.5 - O(\frac{1}{n}) + O(\frac{\sqrt{\varepsilon}}{n})$ and P^2 has prior $P^2(\omega = 1) = 0.5 - O(\frac{1}{n}) - O(\frac{\sqrt{\varepsilon}}{n})$; the conditional distributions of each signal, $P^1(s_i | \omega)$ and $P^2(s_i | \omega)$, differ by $O(\frac{\varepsilon}{n})$ in squared Hellinger distance; taking the product of n signals, $P^1(\mathbf{s} | \omega)$ and $P^2(\mathbf{s} | \omega)$ differ by $O(\varepsilon)$. The distinguishing distributions problem asks: given T samples from either P^1 or P^2 , tell which distribution the samples come from. We show that, if we can solve the forecast aggregation problem, namely, ε -approximate $f^*(\mathbf{r}) = \frac{1}{1 + \rho^{n-1} \prod_{i=1}^n \frac{1-r_i}{r_i}}$, then we can estimate ρ with accuracy $O(\frac{\sqrt{\varepsilon}}{n})$, and hence distinguish P^1 and P^2 . But distinguishing P^1 and P^2 requires $\Omega(\frac{1}{d_{\text{H}}^2(P^1, P^2)} \log \frac{1}{\delta}) = \Omega(\frac{1}{\varepsilon} \log \frac{1}{\delta})$ samples. This gives the lower bound we claimed. See details in Appendix F.1.

5.2 Strongly and Weakly Informative Experts

While the sample complexity of forecast aggregation for general conditionally independent distributions is $O(\frac{1}{\varepsilon^2} \log \frac{1}{\delta})$, under further assumptions this bound can be improved. In particular, we find two special yet natural families of conditionally independent distributions that admit $O(\frac{1}{\varepsilon} \log \frac{1}{\delta})$ sample complexity. In these two cases, the experts are either “very informative” or “very non-informative”. Roughly speaking, an expert is “very informative” if the conditional distributions of the expert’s signal under event $\omega = 0$ and $\omega = 1$ are significantly different, so the expert’s prediction r_i is away from the prior p . An expert is “very non-informative” if the opposite is true. Intuitively, an expert being informative should help aggregation and hence reduce the sample complexity. Interestingly though, we show that even if the experts are non-informative the sample complexity of forecast aggregation can also be reduced. See details in Appendix B.

6 Extension: Multi-Outcome Events

Our main results regarding the sample complexity of forecast aggregation for binary events (Theorems 4.1 and 5.2) can be generalized to multi-outcome events with $|\Omega| > 2$. We prove that: for general distributions, the sample complexity is $\tilde{\Omega}(\frac{m^{n-2}}{\varepsilon}) = \Omega(\frac{m^{n-2} + \log(1/\delta)}{\varepsilon}) \leq T_{\mathcal{P}}(\varepsilon, \delta) \leq O(\frac{|\Omega| m^n + \log(1/\delta)}{\varepsilon^2}) = \tilde{O}(\frac{|\Omega| m^n}{\varepsilon^2})$; for conditionally independent distributions, the sample complexity is $\tilde{\Omega}(\frac{1}{\varepsilon}) = \Omega(\frac{1}{\varepsilon} \log \frac{1}{\delta}) \leq T_{\mathcal{P}_{\text{ind}}}(\varepsilon, \delta) \leq O(\frac{|\Omega| \log |\Omega|}{\varepsilon^2} \log \frac{1}{\varepsilon} + \frac{1}{\varepsilon^2} \log \frac{1}{\delta}) = \tilde{O}(\frac{|\Omega|}{\varepsilon^2})$. See Appendix C for details.

7 Conclusion and Discussion

In this work, we showed an exponential gap between the sample complexity of forecast aggregation for general distributions, $\tilde{\Omega}(\frac{m^{n-2}}{\varepsilon})$, and conditionally independent distributions, $\tilde{O}(\frac{1}{\varepsilon^2})$. This gap is due to the need of estimating the conditional correlation between experts in the general case, which is not needed in the conditional independence case. Notably, the bound $\tilde{O}(\frac{1}{\varepsilon^2})$ for conditionally independent distributions does not depend on the number of experts.

We discuss the dependency of the sample complexity on ε and other directions for future works:

The dependency on ε An open question left by our work is the dependency of the sample complexity on the parameter ε . We conjecture that the tight dependency should be $\frac{1}{\varepsilon}$ (so our lower bounds are tight). This is supported by the following evidence:

Theorem 7.1. *For the case of $|\Omega| = 2$ and for general distributions, if the distribution P has a minimum joint probability $\min_{(\mathbf{s}, \omega) \in \mathcal{S} \times \Omega} P(\mathbf{s}, \omega) > \frac{c}{m^n}$ for some $c > 0$, then the sample complexity of forecast aggregation is at most $O(\frac{m^n}{c\varepsilon} (n \log m + \log \frac{1}{\delta})) = \tilde{O}(\frac{nm^n}{c\varepsilon})$.⁵*

In particular, this theorem can be applied to distributions that are close to uniform, where $P(\mathbf{s}, \omega) \approx \frac{1}{|\mathcal{S} \times \Omega|} = \frac{1}{2m^n}$, giving a bound of $\tilde{O}(\frac{nm^n}{\varepsilon})$. Notably, the set of distributions we constructed in the proof of the $\tilde{\Omega}(\frac{m^{n-2}}{\varepsilon})$ lower bound in Theorem 4.1 are also close to uniform. This means that

⁵This bound has a better dependency on ε but worse on n than the $O(\frac{m^n}{\varepsilon^2})$ bound in Theorem 4.1.

close-to-uniform distributions have a tight sample complexity bound of the form $\Theta(\frac{f(n,m)}{\epsilon})$, not $\Theta(\frac{f(n,m)}{\epsilon^2})$. Moreover, since close-to-uniform distributions are the “most difficult” distributions to learn in the distribution learning problem, it is likely that they are also the most difficult distributions for the forecast aggregation problem, and therefore the tight sample complexity of forecast aggregation should be determined by the sample complexity for those distributions, which is $\Theta(\frac{f(n,m)}{\epsilon})$.

Other future directions

- *The middle ground between fully correlated experts and conditionally independent experts:* An example is the *partial evidence model* in [4]. Applying [4]’s results, one can show that the sample complexity of forecast aggregation in the partial evidence model is at most $\tilde{O}(\frac{n^2}{\epsilon^2})$.⁶ Giving a lower bound for the partial evidence model and exploring other intermediate models is open.
- *Weaker benchmark:* Since the Bayesian aggregator needs exponentially many samples to approximate, can we find a weaker yet meaningful benchmark with a small sample complexity?
- *Samples vs experts:* In reality, obtaining samples of experts’ historical forecasts can be difficult, while recruiting experts is easy. Can we achieve better aggregation by recruiting more experts instead of collecting more samples? How many experts do we need?
- *Eliciting more information:* Previous works on information elicitation and aggregation have noticed that better aggregation can be achieved by eliciting more information than agents’ own predictions, for example, also eliciting each agent’s prediction about other agents’ predictions (e.g., [41, 47, 31, 14]). One can ask whether and how eliciting more information can help to reduce the sample complexity of information aggregation.
- *Continuous distributions:* In our model the random variable ω to be predicted is discrete. One can study a setting where ω is a continuous random variable and the experts report, e.g., the means of their posterior beliefs about ω . The results for continuous random variables might be very different from the results in this work.
- *Other loss functions:* We focused on the squared loss $\mathbb{E}[|f(\mathbf{r}) - \omega|^2]$ due to its popularity in machine learning problems and its useful property that the difference between the squared losses of any aggregator and the optimal aggregator is equal to their expected squared difference (Lemma 2.1). Alternatively, one can consider other loss functions like the logarithmic loss $\mathbb{E}[\omega \log(f(\mathbf{r})) + (1 - \omega) \log(1 - f(\mathbf{r}))]$ and the absolute loss $\mathbb{E}[|f(\mathbf{r}) - \omega|]$. There might be some technical challenges in the analysis of sample complexity for those loss functions, though: e.g., the logarithmic loss can be unbounded [4, 40] and the absolute loss does not enjoy a property like Lemma 2.1.

Acknowledgments and Disclosure of Funding

We would like to thank Yannai Gonczarowski, Ariel Procaccia, Milind Tambe, David Parkes, Bo Waggoner, Rafael Frongillo, Grant Schoenebeck, Yuqing Kong, and anonymous reviewers for their helpful comments. This research is based upon work supported in part by the National Science Foundation under grants no. IIS-2007887 and no. IIS-2147187.

References

- [1] Martin Anthony and Peter L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, 1 edition, November 1999.
- [2] Itai Arieli, Yakov Babichenko, and Rann Smorodinsky. Robust forecast aggregation. *Proceedings of the National Academy of Sciences*, 115(52), December 2018.
- [3] Itai Arieli, Yakov Babichenko, Inbal Talgam-Cohen, and Konstantin Zabarnyi. Universally Robust Information Aggregation for Binary Decisions. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pages 118–118, London United Kingdom, July 2023. ACM.
- [4] Yakov Babichenko and Dan Garber. Learning Optimal Forecast Aggregation in Partial Evidence Environments. *Mathematics of Operations Research*, 46(2):628–641, May 2021.

⁶In particular, [4] studies an online learning setting with logarithmic loss. Their regret bound can be converted to our sample complexity bound by an online-to-batch conversion and the Pinsker’s inequality.

- [5] Maria-Florina F Balcan, Tuomas Sandholm, and Ellen Vitercik. Sample complexity of automated mechanism design. *Advances in Neural Information Processing Systems*, 29, 2016.
- [6] Jonathan Baron, Barbara A. Mellers, Philip E. Tetlock, Eric Stone, and Lyle H. Ungar. Two Reasons to Make Aggregated Probability Forecasts More Extreme. *Decision Analysis*, 11(2):133–145, June 2014.
- [7] John M. Bates and Clive W. J. Granger. The Combination of Forecasts. *Journal of the Operational Research Society*, 20(4):451–468, December 1969.
- [8] Robert F. Bordley. A Multiplicative Formula for Aggregating Probability Assessments. *Management Science*, 28(10):1137–1148, 1982. Publisher: INFORMS.
- [9] Glenn W. Brier. Verification of Forecasts Expressed in Terms of Probability. *Monthly Weather Review*, 78(1):1–3, January 1950.
- [10] Johannes Brustle, Yang Cai, and Constantinos Daskalakis. Multi-Item Mechanisms without Item-Independence: Learnability via Robustness. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 715–761, Virtual Event Hungary, July 2020. ACM.
- [11] Yang Cai and Constantinos Daskalakis. Learning Multi-Item Auctions with (or without) Samples. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 516–527, Berkeley, CA, October 2017. IEEE.
- [12] Clément L. Canonne. A short note on learning discrete distributions, 2020.
- [13] Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, Cambridge, 2006.
- [14] Yi-Chun Chen, Manuel Mueller-Frank, and Mallesh Pai. The Wisdom of the Crowd and Higher-Order Beliefs. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pages 450–450, London United Kingdom, July 2023. ACM.
- [15] Gerda Claeskens, Jan R. Magnus, Andrey L. Vasnev, and Wendun Wang. The forecast combination puzzle: A simple theoretical explanation. *International Journal of Forecasting*, 32(3):754–762, July 2016.
- [16] Robert T. Clemen. Combining forecasts: A review and annotated bibliography. *International Journal of Forecasting*, 5(4):559–583, January 1989.
- [17] Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing - STOC '14*, pages 243–252, New York, New York, 2014. ACM Press.
- [18] Henrique De Oliveira, Yuhta Ishii, and Xiao Lin. Robust Merging of Information. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 341–342, Budapest Hungary, July 2021. ACM.
- [19] Nikhil R. Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. The sample complexity of auctions with side information. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 426–439, Cambridge MA USA, June 2016. ACM.
- [20] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333, May 2015.
- [21] R. Glen Donaldson and Mark Kamstra. Forecast combining with neural networks. *Journal of Forecasting*, 15(1):49–61, January 1996.
- [22] Rafael Frongillo, Robert Gomez, Anish Thilagar, and Bo Waggoner. Efficient Competitions and Online Learning with Strategic Forecasters. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 479–496, Budapest Hungary, July 2021. ACM.
- [23] Véronique Genre, Geoff Kenny, Aidan Meyler, and Allan Timmermann. Combining expert forecasts: Can anything beat the simple average? *International Journal of Forecasting*, 29(1):108–121, January 2013.
- [24] John Geweke and Charles Whiteman. Chapter 1 Bayesian Forecasting. In *Handbook of Economic Forecasting*, volume 1, pages 3–80. Elsevier, 2006.
- [25] Yannai A. Gonczarowski and S. Matthew Weinberg. The Sample Complexity of Up-to- ϵ Multi-dimensional Revenue Maximization. *Journal of the ACM*, 68(3):1–28, March 2021.

- [26] Clive W. J. Granger. Invited review combining forecasts—twenty years later. *Journal of Forecasting*, 8(3):167–173, July 1989.
- [27] Clive W. J. Granger and Ramu Ramanathan. Improved methods of combining forecasts. *Journal of Forecasting*, 3(2):197–204, April 1984.
- [28] Chenghao Guo, Zhiyi Huang, Zhihao Gavin Tang, and Xinzhi Zhang. Generalizing complex hypotheses on product distributions: Auctions, prophet inequalities, and pandora’s problem. In *Conference on Learning Theory*, pages 2248–2288. PMLR, 2021.
- [29] Chenghao Guo, Zhiyi Huang, and Xinzhi Zhang. Settling the sample complexity of single-parameter revenue maximization. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing - STOC 2019*, pages 662–673, Phoenix, AZ, USA, 2019. ACM Press.
- [30] Sunil Gupta and Peter C. Wilton. Combination of Forecasts: An Extension. *Management Science*, 33(3):356–372, 1987. Publisher: INFORMS.
- [31] Yuqing Kong, Yunqi Li, Yubo Zhang, Zhihuan Huang, and Jinzhao Wu. Eliciting Thinking Hierarchy without a Prior. In *Advances in Neural Information Processing Systems*, volume 35, pages 13329–13341. Curran Associates, Inc., 2022.
- [32] Jasper Lee. Lecture 11: Distinguishing (discrete) distributions, 2020.
- [33] Jasper Lee. Lecture 12: Learning Discrete Distributions, 2020.
- [34] Spyros Makridakis, Evangelos Spiliotis, and Vassilios Assimakopoulos. The M4 Competition: 100,000 time series and 61 forecasting methods. *International Journal of Forecasting*, 36(1):54–74, January 2020.
- [35] Spyros Makridakis and Robert L. Winkler. Averages of Forecasts: Some Empirical Results. *Management Science*, 29(9):987–996, September 1983.
- [36] Jamie H Morgenstern and Tim Roughgarden. On the pseudo-dimension of nearly optimal auctions. *Advances in Neural Information Processing Systems*, 28, 2015.
- [37] Peter A. Morris. Decision Analysis Expert Use. *Management Science*, 20(9):1233–1241, 1974. Publisher: INFORMS.
- [38] Paul Newbold and Clive W. J. Granger. Experience with Forecasting Univariate Time Series and the Combination of Forecasts. *Journal of the Royal Statistical Society. Series A (General)*, 137(2):131–146, 1974.
- [39] Eric Neyman and Tim Roughgarden. Are You Smarter Than a Random Expert? The Robust Aggregation of Substitutable Signals. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 990–1012, Boulder CO USA, July 2022. ACM.
- [40] Eric Neyman and Tim Roughgarden. No-regret learning with unbounded losses: The case of logarithmic pooling. *arXiv preprint arXiv:2202.11219*, 2022.
- [41] Dražen Prelec, H. Sebastian Seung, and John McCoy. A solution to the single-question crowd wisdom problem. *Nature*, 541(7638):532–535, January 2017.
- [42] Ville A. Satopää, Jonathan Baron, Dean P. Foster, Barbara A. Mellers, Philip E. Tetlock, and Lyle H. Ungar. Combining multiple probability predictions using a simple logit model. *International Journal of Forecasting*, 30(2):344–356, April 2014.
- [43] Jeremy Smith and Kenneth F. Wallis. A Simple Explanation of the Forecast Combination Puzzle. *Oxford Bulletin of Economics and Statistics*, 71(3):331–355, June 2009.
- [44] James H. Stock and Mark W. Watson. Combination forecasts of output growth in a seven-country data set. *Journal of Forecasting*, 23(6):405–430, September 2004.
- [45] Allan Timmermann. Chapter 4 Forecast Combinations. In *Handbook of Economic Forecasting*, volume 1, pages 135–196. Elsevier, 2006.
- [46] Alexandre B. Tsybakov. *Introduction to nonparametric estimation*. Springer series in statistics. Springer, New York ; London, 2009. OCLC: ocn300399286.
- [47] Juntao Wang, Yang Liu, and Yiling Chen. Forecast Aggregation via Peer Prediction. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, 9(1):131–142, October 2021.

- [48] Robert L. Winkler. Combining Probability Distributions from Dependent Information Sources. *Management Science*, 27(4):479–488, 1981. Publisher: INFORMS.
- [49] Robert L. Winkler and Spyros Makridakis. The Combination of Forecasts. *Journal of the Royal Statistical Society. Series A (General)*, 146(2):150, 1983.
- [50] Chunxue Yang and Xiaohui Bei. Learning Optimal Auctions with Correlated Valuations from Samples. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 11716–11726. PMLR, July 2021.
- [51] Yuhong Yang. Combining Forecasting Procedures: Some Theoretical Results. *Econometric Theory*, 20(1):176–222, 2004. Publisher: Cambridge University Press.

A Additional Preliminaries

A.1 Properties of Hellinger Distance

We review some useful properties of the Hellinger distance. First, the Hellinger distance gives upper bounds on the total variation distance:

Fact A.1 (e.g., Lemma 2.3 in [46]).

- $d_{\text{TV}}(D_1, D_2) \leq \sqrt{2}d_{\text{H}}(D_1, D_2)$.
- $1 - d_{\text{TV}}^2(D_1, D_2) \geq (1 - d_{\text{H}}^2(D_1, D_2))^2$. (This inequality implies the first one.)

Second, we will use the following lemma to upper bound the squared Hellinger distance between two distributions that are close to each other:

Lemma A.2. *Let D_1 and D_2 be two distributions on \mathcal{X} satisfying $1 - \varepsilon \leq \frac{D_2(x)}{D_1(x)} \leq 1 + \varepsilon$, $\forall x \in \mathcal{X}$. Then, $d_{\text{H}}^2(D_1, D_2) \leq \frac{1}{2}\varepsilon^2$.*

Proof. By definition,

$$d_{\text{H}}^2(D_1, D_2) = \frac{1}{2} \sum_{x \in \mathcal{X}} \left(\sqrt{D_1(x)} - \sqrt{D_2(x)} \right)^2 = \frac{1}{2} \sum_{x \in \mathcal{X}} D_1(x) \left(1 - \sqrt{\frac{D_2(x)}{D_1(x)}} \right)^2.$$

If $1 - \varepsilon \leq \frac{D_2(x)}{D_1(x)} < 1$, then we have $\left(1 - \sqrt{\frac{D_2(x)}{D_1(x)}} \right)^2 \leq (1 - \sqrt{1 - \varepsilon})^2 \leq (1 - (1 - \varepsilon))^2 = \varepsilon^2$. If $1 + \varepsilon \geq \frac{D_2(x)}{D_1(x)} > 1$, then we have $\left(1 - \sqrt{\frac{D_2(x)}{D_1(x)}} \right)^2 \leq (\sqrt{1 + \varepsilon} - 1)^2 \leq ((1 + \varepsilon) - 1)^2 = \varepsilon^2$. These two cases together imply

$$d_{\text{H}}^2(D_1, D_2) \leq \frac{1}{2} \sum_{x \in \mathcal{X}} D_1(x) \cdot \varepsilon^2 = \frac{1}{2}\varepsilon^2. \quad \square$$

Third, we will use a property of the Hellinger distance between distributions defined on a product space. Suppose D_1 and D_2 are two distributions over a product space $\mathcal{X} \times \mathcal{Y}$. They can be decomposed into the marginal distribution of $x \in \mathcal{X}$ and the conditional distribution of $y \in \mathcal{Y}$ given x , namely, $D_1(x, y) = D_{1,x}(x) \cdot D_{1,y|x}(y|x)$ and $D_2(x, y) = D_{2,x}(x) \cdot D_{2,y|x}(y|x)$. Then, the squared Hellinger distance between D_1 and D_2 satisfies the following:

Lemma A.3. $d_{\text{H}}^2(D_1, D_2) \leq d_{\text{H}}^2(D_{1,x}, D_{2,x}) + \max_{x \in \mathcal{X}} d_{\text{H}}^2(D_{1,y|x}, D_{2,y|x})$. *In particular, if x and y are independent, then $d_{\text{H}}^2(D_1, D_2) \leq d_{\text{H}}^2(D_{1,x}, D_{2,x}) + d_{\text{H}}^2(D_{1,y}, D_{2,y})$.*

Proof. By definition,

$$\begin{aligned} d_{\text{H}}^2(D_1, D_2) &= 1 - \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} \sqrt{D_1(x, y) D_2(x, y)} \\ &= 1 - \sum_{x \in \mathcal{X}} \sqrt{D_{1,x}(x) D_{2,x}(x)} \sum_{y \in \mathcal{Y}} \sqrt{D_{1,y|x}(y|x) D_{2,y|x}(y|x)} \\ &= 1 - \sum_{x \in \mathcal{X}} \sqrt{D_{1,x}(x) D_{2,x}(x)} + \sum_{x \in \mathcal{X}} \sqrt{D_{1,x}(x) D_{2,x}(x)} \left(1 - \sum_{y \in \mathcal{Y}} \sqrt{D_{1,y|x}(y|x) D_{2,y|x}(y|x)} \right) \\ &= d_{\text{H}}^2(D_{1,x}, D_{2,x}) + \sum_{x \in \mathcal{X}} \sqrt{D_{1,x}(x) D_{2,x}(x)} \cdot d_{\text{H}}^2(D_{1,y|x}, D_{2,y|x}) \\ &\leq d_{\text{H}}^2(D_{1,x}, D_{2,x}) + \sum_{x \in \mathcal{X}} \sqrt{D_{1,x}(x) D_{2,x}(x)} \cdot \max_{x \in \mathcal{X}} d_{\text{H}}^2(D_{1,y|x}, D_{2,y|x}) \\ &\leq d_{\text{H}}^2(D_{1,x}, D_{2,x}) + 1 \cdot \max_{x \in \mathcal{X}} d_{\text{H}}^2(D_{1,y|x}, D_{2,y|x}), \end{aligned}$$

where the last inequality is because $\sum_{x \in \mathcal{X}} \sqrt{D_{1,x}(x) D_{2,x}(x)} = 1 - d_{\text{H}}^2(D_{1,x}, D_{2,x}) \leq 1$. \square

Finally, let $D^{\otimes T}$ denote the distribution of T i.i.d. samples from D , namely, the product of T independent D 's. We have the following lemma that relates $d_{\text{H}}^2(D_1^{\otimes T}, D_2^{\otimes T})$ with $d_{\text{H}}^2(D_1, D_2)$:

Lemma A.4 (e.g., [32]). $d_{\text{H}}^2(D_1^{\otimes T}, D_2^{\otimes T}) = 1 - (1 - d_{\text{H}}^2(D_1, D_2))^T \leq T \cdot d_{\text{H}}^2(D_1, D_2)$.

A.2 Distinguishing Distributions

Let D_1, D_2 be two distributions over a discrete space \mathcal{X} . A distribution D_i is chosen uniformly at random from $\{D_1, D_2\}$. Then, we are given T samples from D_i and want to guess whether the distribution is D_1 or D_2 . It is known that at least $T = \Omega\left(\frac{1}{d_{\text{H}}^2(D_1, D_2)} \log \frac{1}{\delta}\right)$ samples are needed to guess correctly with probability at least $1 - \delta$, no matter how we guess. Formally:

Lemma A.5 (e.g., [32]). *Let $j \in \{1, 2\}$ be the index of the distribution we guess based on the samples. The probability of making a mistake when distinguishing D_1 and D_2 using T samples, namely $\Pr[j \neq i] = \frac{1}{2} \Pr[j \neq i | i = 1] + \frac{1}{2} \Pr[j \neq i | i = 2]$, is at least:*

- $\Pr[j \neq i] \geq \frac{1}{2} - \sqrt{\frac{T}{2}} d_{\text{H}}(D_1, D_2)$.
- $\Pr[j \neq i] \geq \frac{1}{4} (1 - d_{\text{H}}^2(D_1, D_2))^{2T} \geq \frac{1}{4} e^{-4T d_{\text{H}}^2(D_1, D_2)}$, assuming $d_{\text{H}}^2(D_1, D_2) \leq \frac{1}{2}$.

The second item implies that, in order to achieve $\Pr[j \neq i] \leq \delta$, we must have $T \geq \frac{1}{4d_{\text{H}}^2(D_1, D_2)} \log \frac{1}{4\delta}$.

We provide a proof of this lemma for completeness:

Proof. Let $D_1^{\otimes T}$ and $D_2^{\otimes T}$ denote the distributions of T i.i.d. samples from D_1 and D_2 , respectively. The draw of T samples from D_1 or D_2 is equivalent to the draw of one sample from $D_1^{\otimes T}$ or $D_2^{\otimes T}$. Given one sample from $D_1^{\otimes T}$ or $D_2^{\otimes T}$, the probability of making a mistake when guessing the distribution is at least:

$$\begin{aligned}
\Pr[j \neq i] &= \frac{1}{2} \left(\Pr[j = 2 | i = 1] + \Pr[j = 1 | i = 2] \right) \\
&= \frac{1}{2} \left(1 - \Pr[j = 1 | i = 1] + \Pr[j = 1 | i = 2] \right) \\
&= \frac{1}{2} - \frac{1}{2} \left(\Pr[j = 1 | i = 1] - \Pr[j = 1 | i = 2] \right) \\
&= \frac{1}{2} - \frac{1}{2} \left(\mathbb{E}_{D_1^{\otimes T}}[\mathbb{1}\{j = 1\}] - \mathbb{E}_{D_2^{\otimes T}}[\mathbb{1}\{j = 1\}] \right) \\
&\text{by Fact 3.2} \geq \frac{1}{2} - \frac{1}{2} d_{\text{TV}}(D_1^{\otimes T}, D_2^{\otimes T}). \tag{12}
\end{aligned}$$

We then upper bound $d_{\text{TV}}(D_1^{\otimes T}, D_2^{\otimes T})$ in two ways, which will prove the two items of the lemma, respectively. First, according to first item of Fact A.1, we have

$$d_{\text{TV}}(D_1^{\otimes T}, D_2^{\otimes T}) \leq \sqrt{2} d_{\text{H}}(D_1^{\otimes T}, D_2^{\otimes T}).$$

By Lemma A.4,

$$d_{\text{H}}(D_1^{\otimes T}, D_2^{\otimes T}) \leq \sqrt{T} d_{\text{H}}(D_1, D_2).$$

The above two inequalities give $d_{\text{TV}}(D_1^{\otimes T}, D_2^{\otimes T}) \leq \sqrt{2T} d_{\text{H}}(D_1, D_2)$. This proves the first item of the lemma.

Second, according to the second item of Fact A.1 and Lemma A.4, we have

$$1 - d_{\text{TV}}^2(D_1^{\otimes T}, D_2^{\otimes T}) \geq (1 - d_{\text{H}}^2(D_1^{\otimes T}, D_2^{\otimes T}))^2 = (1 - d_{\text{H}}^2(D_1, D_2))^{2T}$$

Since $1 - d_{\text{TV}}^2(D_1^{\otimes T}, D_2^{\otimes T}) = (1 + d_{\text{TV}}(D_1^{\otimes T}, D_2^{\otimes T}))(1 - d_{\text{TV}}(D_1^{\otimes T}, D_2^{\otimes T})) \leq 2(1 - d_{\text{TV}}(D_1^{\otimes T}, D_2^{\otimes T}))$, we have

$$1 - d_{\text{TV}}(D_1^{\otimes T}, D_2^{\otimes T}) \geq \frac{1}{2} (1 - d_{\text{H}}^2(D_1, D_2))^{2T}.$$

Plugging into (12), we get

$$\Pr[j \neq i] \geq \frac{1}{4} (1 - d_{\text{H}}^2(D_1, D_2))^{2T}.$$

When $d_{\text{H}}^2(D_1, D_2) < \frac{1}{2}$, we use the inequality $1 - x \geq e^{-2x}$ for $0 < x < \frac{1}{2}$ to conclude that

$$\Pr[j \neq i] \geq \frac{1}{4} (e^{-2d_{\text{H}}^2(D_1, D_2)})^{2T} = \frac{1}{4} e^{-4Td_{\text{H}}^2(D_1, D_2)}. \quad \square$$

B Special Cases: Strongly and Weakly Informative Experts

In this section we investigate two special families of conditionally independent distributions that admit $O(\frac{1}{\varepsilon} \log \frac{1}{\delta})$ sample complexity, which is smaller than the $O(\frac{1}{\varepsilon^2} \log \frac{1}{\varepsilon\delta})$ bound for general conditionally independent distributions. In these two cases, the (signals of) experts are either “very informative” or “very non-informative”.

Definition B.1. Let $\gamma \in [0, \infty]$ be a parameter. For an expert i , we say its signal $s_i \in \mathcal{S}_i$ is

- γ -strongly informative if either $\frac{P(s_i|\omega=1)}{P(s_i|\omega=0)} \geq 1 + \gamma$ or $\frac{P(s_i|\omega=1)}{P(s_i|\omega=0)} \leq \frac{1}{1+\gamma}$ holds.
- γ -weakly informative if $\frac{1}{1+\gamma} \leq \frac{P(s_i|\omega=1)}{P(s_i|\omega=0)} \leq 1 + \gamma$.

An expert i is γ -strongly (or γ -weakly) informative if all of its signals in \mathcal{S}_i are γ -strongly (or γ -weakly) informative.⁷

A signal s_i being γ -strongly (or γ -weakly) informative implies that its corresponding report r_i will be “ γ -away from” (or “ γ -close to”) the prior $p = P(\omega = 1)$, in the “ $\frac{r_i}{1-r_i}$ and $\frac{p}{1-p}$ ” ratio form. Specifically, if s_i is γ -strongly informative, then from Equation (1) we have

$$\frac{r_i}{1-r_i} = \frac{P(s_i|\omega=1)}{P(s_i|\omega=0)} \frac{p}{1-p} \geq (1 + \gamma)\rho \quad \text{or} \quad \leq \frac{1}{1+\gamma}\rho. \quad (13)$$

As γ gets larger, a γ -strongly informative signal (expert) is *more* informative for predicting whether $\omega = 1$ or 0. This would make aggregation easier. If s_i is γ -weakly informative, then:

$$\frac{1}{1+\gamma}\rho \leq \frac{r_i}{1-r_i} \leq (1 + \gamma)\rho. \quad (14)$$

As γ gets smaller, a γ -weakly informative signal (expert) is *less* informative for predicting ω , but in this case their report r_i will be close to the prior p , which allows us to estimate the ρ^{n-1} term in the optimal aggregator $f^*(\mathbf{r}) = \frac{1}{1+\rho^{n-1} \prod_{i=1}^n \frac{1-r_i}{r_i}}$ better. Those are some intuitions why both strongly and weakly informative signals can reduce the sample complexity of forecast aggregation.

Formally, for γ -strongly informative experts with not-too-small γ , we have the following result:

Theorem B.2. If $n \geq 32 \log \frac{2}{\varepsilon}$ and all experts are γ -strongly informative with $\frac{\gamma}{1+\gamma} \geq 8\sqrt{\frac{2}{n} \log \frac{2}{\varepsilon}}$, then the sample complexity of forecast aggregation is $\leq O(\frac{1}{\varepsilon n (\frac{\gamma}{1+\gamma})^2} \log \frac{1}{\delta} + \frac{1}{\varepsilon} \log \frac{1}{\delta}) = O(\frac{1}{\varepsilon} \log \frac{1}{\delta})$.

We remark that the conditions of the theorem, $n \geq 32 \log \frac{2}{\varepsilon}$ and $\frac{\gamma}{1+\gamma} \geq 8\sqrt{\frac{2}{n} \log \frac{2}{\varepsilon}}$, are easier to be satisfied when the number of experts n increases, if the informativeness parameter γ of each expert is a constant or does not decrease with n (which we believe is a reasonable assumption given that experts are independent of each other). Also, if γ is fixed or increasing, then as n increases the sample complexity decreases.

The proof of Theorem B.2 is in Appendix G.1. Roughly speaking, we divide each expert’s signal set into two sets, \mathcal{S}_i^1 and \mathcal{S}_i^0 : signals that are more likely to occur under $\omega = 1$ (i.e., $\frac{P(s_i|\omega=1)}{P(s_i|\omega=0)} \geq 1 + \gamma$) and under $\omega = 0$ (i.e., $\frac{P(s_i|\omega=1)}{P(s_i|\omega=0)} \leq \frac{1}{1+\gamma}$). If the realized ω is 1, then one may expect to see $\Omega(\frac{\gamma}{1+\gamma}n)$ more \mathcal{S}_i^1 signals than the \mathcal{S}_i^0 signals from the n experts, because the probabilities of these two types of signals differ by $\Omega(\frac{\gamma}{1+\gamma})$ for each expert. If ω is 0 then one may expect to see more \mathcal{S}_i^0 signals than the \mathcal{S}_i^1 signals. So, by checking which type of signals are more we can tell whether $\omega = 0$ or

⁷We note that an expert can be neither γ -strongly informative nor γ -weakly informative for any γ .

1. To tell whether a signal belongs to \mathcal{S}_i^1 or \mathcal{S}_i^0 , we compare the corresponding report $\frac{r_i}{1-r_i}$ with ρ (namely, $\frac{r_i}{1-r_i} \geq (1+\gamma)\rho \iff s_i \in \mathcal{S}_i^1$ and $\frac{r_i}{1-r_i} \leq \frac{1}{1+\gamma}\rho \iff s_i \in \mathcal{S}_i^0$), where ρ is estimated with accuracy $\frac{\gamma}{1+\gamma}$ using $O(\frac{1}{\rho n (\frac{\gamma}{1+\gamma})^2} \log \frac{1}{\delta})$ samples. Dealing with the case of $\rho < \varepsilon$ separately, we obtain the bound $O(\frac{1}{\varepsilon n (\frac{\gamma}{1+\gamma})^2} \log \frac{1}{\delta} + \frac{1}{\varepsilon} \log \frac{1}{\delta})$.

For γ -weakly informative experts with small γ , we have:

Theorem B.3. *If all experts are γ -weakly informative with $\gamma \leq 1$, then the sample complexity of forecast aggregation is $\leq \min \{O(\frac{\gamma n}{\varepsilon} \log \frac{1}{\delta}), O(\frac{1}{\varepsilon^2} \log \frac{1}{\delta})\}$, which is $O(\frac{1}{\varepsilon} \log \frac{1}{\delta})$ if $\gamma = O(\frac{1}{n})$.*

The $O(\frac{1}{\varepsilon^2} \log \frac{1}{\delta})$ term in the sample complexity follows from the result for general conditionally independent distributions (Theorem 5.2). The $O(\frac{\gamma n}{\varepsilon} \log \frac{1}{\delta})$ term is proved in Appendix G.2. We give the rough idea here using $\gamma = \frac{1}{n}$ as an example. The proof relies on the observation that $\mathbb{E}[\prod_{i=1}^n \frac{r_i}{1-r_i} \mid \omega = 0] = \rho^n$. Since experts are weakly informative, each of their reports $\frac{r_i}{1-r_i}$ is around ρ in the range $[\frac{1}{1+\gamma}\rho, (1+\gamma)\rho] \subseteq [\frac{1}{1+\frac{1}{n}}\rho, (1+\frac{1}{n})\rho]$. Taking the product, we have $\prod_{i=1}^n \frac{r_i}{1-r_i} \in [\frac{\rho^n}{e}, e\rho^n]$, which is in a bounded range. This allows us to use Chernoff bound to argue that the ρ^n (or ρ^{n-1}) term in the optimal aggregator $f^*(\mathbf{r}) = \frac{1}{1+\rho^{n-1} \prod_{i=1}^n \frac{1-r_i}{r_i}}$ can be estimated, with $O(\sqrt{\varepsilon})$ accuracy, using only $O(\frac{e}{(\sqrt{\varepsilon})^2} \log \frac{1}{\delta}) = O(\frac{1}{\varepsilon} \log \frac{1}{\delta})$ samples of $\prod_{i=1}^n \frac{r_i}{1-r_i}$. The aggregator using the estimate $\hat{\rho}^{n-1}$, $\hat{f}(\mathbf{r}) = \frac{1}{1+\hat{\rho}^{n-1} \prod_{i=1}^n \frac{1-r_i}{r_i}}$, turns out to be $O(\varepsilon)$ -optimal.

C Extension: Multi-Outcome Events

In this section we generalize our model from binary events to multi-outcome events. The event space now becomes $\Omega = \{1, 2, \dots, |\Omega|\}$ with $|\Omega| > 2$. The joint distribution of event $\omega \in \Omega$ and experts' signals $\mathbf{s} = (s_i)_{i=1}^n \in \mathcal{S}$ is still denoted by P , which belongs to some class of distributions \mathcal{P} . The size of each expert's signal space is still assumed to be $|\mathcal{S}_i| = m < +\infty$ for general distributions and can be infinite for conditionally independent distributions (where s_1, \dots, s_n are independent conditioned on ω). After observing signal s_i , expert i reports its posterior belief of the event given s_i , which is $r_i = (r_{ij})_{j \in \Omega}$ where $r_{ij} = P(\omega = j \mid s_i)$. An aggregator now is a vector-valued function $\mathbf{f} = (f_j)_{j \in \Omega}$ that maps the joint report $\mathbf{r} = (\mathbf{r}_i)_{i=1}^n = (r_{ij})_{ij}$ to a probability distribution $\mathbf{f}(\mathbf{r})$ over Ω , where $f_j(\mathbf{r})$ is the aggregated predicted probability for $\omega = j$. We assume $f_j(\mathbf{r}) \geq 0$ and $\sum_{j \in \Omega} f_j(\mathbf{r}) = 1$. The definition of the (expected) loss of an aggregator \mathbf{f} becomes:

$$L_P(\mathbf{f}) = \mathbb{E}_P \left[\sum_{j \in \Omega} |f_j(\mathbf{r}) - \mathbb{1}[\omega = j]|^2 \right]. \quad (15)$$

It is easy to see that the optimal aggregator $\mathbf{f}^* = \operatorname{argmin}_{\mathbf{f}} L_P(\mathbf{f})$ in the multi-outcome case is still the Bayesian aggregator (this is a generalization of Lemma 2.1):

$$\mathbf{f}^* = (f_j^*)_{j \in \Omega}, \quad f_j^*(\mathbf{r}) = P(\omega = j \mid \mathbf{r}) \quad (16)$$

and the difference between the losses of \mathbf{f} and \mathbf{f}^* satisfies

$$L_P(\mathbf{f}) - L_P(\mathbf{f}^*) = \mathbb{E}_P \left[\sum_{j \in \Omega} |f_j(\mathbf{r}) - f_j^*(\mathbf{r})|^2 \right]. \quad (17)$$

So, an ε -optimal aggregator \mathbf{f} is an aggregator that satisfies $\mathbb{E}_P [\sum_{j \in \Omega} |f_j(\mathbf{r}) - f_j^*(\mathbf{r})|^2] \leq \varepsilon$.

Using $T = T_{\mathcal{P}}(\varepsilon, \delta)$ samples $\{(\mathbf{r}^{(t)}, \omega^{(t)})\}_{t=1}^T$ from P , the principal wants to find an ε -optimal aggregator $\hat{\mathbf{f}}$ with probability at least $1 - \delta$, for any distribution P in a class \mathcal{P} . We give lower bounds and upper bounds on the sample complexity:

Theorem C.1. *The sample complexity of forecast aggregation for multi-outcome events is:*

- For the class \mathcal{P} of general distributions: $\Omega \left(\frac{m^{n-2} + \log(1/\delta)}{\varepsilon} \right) \leq T_{\mathcal{P}}(\varepsilon, \delta) \leq O \left(\frac{|\Omega| m^n + \log(1/\delta)}{\varepsilon^2} \right)$.
- For the class \mathcal{P}_{ind} of conditionally independent distributions: $\Omega \left(\frac{1}{\varepsilon} \log \frac{1}{\delta} \right) \leq T_{\mathcal{P}_{\text{ind}}}(\varepsilon, \delta) \leq O \left(\frac{|\Omega| \log |\Omega|}{\varepsilon^2} \log \frac{1}{\varepsilon} + \frac{1}{\varepsilon^2} \log \frac{1}{\delta} \right)$.

C.1 Proof of Theorem C.1

Before proving the theorem, we note that the loss of any aggregator \mathbf{f} satisfying $f_j(\mathbf{r}) \geq 0$ and $\sum_{j \in \Omega} f_j(\mathbf{r}) = 1$ is bounded by $[0, 2]$:

$$0 \leq \sum_{j \in \Omega} |f_j(\mathbf{r}) - \mathbb{1}[\omega = j]|^2 \leq \sum_{j \in \Omega} |f_j(\mathbf{r}) - \mathbb{1}[\omega = j]| \leq 1 + \sum_{j \in \Omega} f_j(\mathbf{r}) = 2. \quad (18)$$

Similarly,

$$0 \leq \sum_{j \in \Omega} |f_j(\mathbf{r}) - f_j^*(\mathbf{r})|^2 \leq 2. \quad (19)$$

C.1.1 Lower Bounds

The lower bounds directly follow from the lower bounds for the binary case (Theorem 4.1 and Theorem 5.2) because the binary case is a special case of the multi-outcome case. Specifically, we can regard any binary event distribution P as a multi-outcome distribution that puts probability only on outcomes $\{1, 2\} \subseteq \Omega$. If we can learn an ε -optimal aggregator $\hat{\mathbf{f}}$ for the multi-outcome case: $\mathbb{E}[\sum_{j \in \Omega} |\hat{f}_j(\mathbf{r}) - f_j^*(\mathbf{r})|^2] \leq \varepsilon$, then this aggregator satisfies

$$\begin{aligned} & \mathbb{E}\left[\sum_{j \in \{1,2\}} |\hat{f}_j(\mathbf{r}) - f_j^*(\mathbf{r})|^2\right] \leq \varepsilon \\ \iff & \mathbb{E}\left[|\hat{f}_1(\mathbf{r}) - f_1^*(\mathbf{r})|^2 + |(1 - \hat{f}_1(\mathbf{r})) - (1 - f_1^*(\mathbf{r}))|^2\right] = \mathbb{E}\left[2|\hat{f}_1(\mathbf{r}) - f_1^*(\mathbf{r})|^2\right] \leq \varepsilon \\ \iff & \mathbb{E}\left[|\hat{f}_1(\mathbf{r}) - f_1^*(\mathbf{r})|^2\right] \leq \frac{\varepsilon}{2}. \end{aligned}$$

So, the aggregator $\hat{f}_1(\cdot)$ is an $\frac{\varepsilon}{2}$ -optimal aggregator for the binary case.

C.1.2 Upper Bounds

General distributions: the $O\left(\frac{|\Omega|m^n + \log(1/\delta)}{\varepsilon^2}\right)$ upper bound. The proof for general distributions in the multi-outcome case is the same as the proof for general distributions in the binary case (Section 4.1), except for three differences: (1) P (regarded as a distribution over reports \mathbf{r} and event ω) now is a discrete distribution with support size at most $|\Omega|m^n$; (2) the loss $|f(\mathbf{r}) - \omega|^2 \in [0, 1]$ in the binary case is replaced by the loss $\sum_{j \in \Omega} |f_j(\mathbf{r}) - \mathbb{1}[\omega = j]|^2 \in [0, 2]$ in the multi-outcome case; (3) the ε in the binary case is replaced by $\frac{\varepsilon}{2}$ because the loss is now upper bounded by 2. Thus, the bound $O\left(\frac{m^n + \log(1/\delta)}{\varepsilon^2}\right)$ in the binary case becomes $O\left(\frac{|\Omega|m^n + \log(1/\delta)}{(\frac{\varepsilon}{2})^2}\right) = O\left(\frac{|\Omega|m^n + \log(1/\delta)}{\varepsilon^2}\right)$ in the multi-outcome case.

Conditionally independent distributions: the $O\left(\frac{|\Omega| \log |\Omega|}{\varepsilon^2} \log \frac{1}{\varepsilon} + \frac{1}{\varepsilon^2} \log \frac{1}{\delta}\right)$ upper bound. The sample complexity upper bound for conditionally independent distributions is proved by a *pseudo-dimension* argument. We first show in Lemma C.2 that the optimal aggregator \mathbf{f}^* belongs to some parametric family of aggregators. Then, we upper bound the pseudo-dimension of the loss functions associated with this parametric family of aggregators, which will give the desired sample complexity upper bound.

Lemma C.2. *For multi-outcome conditionally independent distribution P , given signals $\mathbf{s} = (s_i)_{i=1}^n$, with corresponding reports $\mathbf{r} = (r_i)_{i=1}^n = (r_{ij})_{ij}$ where $r_{ij} = P(\omega = j | s_i)$, the posterior probability of $\omega = j$ satisfies:*

$$f_j^*(\mathbf{r}) = P(\omega = j | \mathbf{r}) = P(\omega = j | \mathbf{s}) = \frac{\frac{1}{P(\omega=j)^{n-1}} \prod_{i=1}^n r_{ij}}{\sum_{k \in \Omega} \frac{1}{P(\omega=k)^{n-1}} \prod_{i=1}^n r_{ik}}.$$

Proof. By Bayes' rule and the fact that s_1, \dots, s_n are independent conditioning on ω ,

$$\begin{aligned}
P(\omega = j \mid \mathbf{s}) &= \frac{P(\omega = j) \prod_{i=1}^n P(s_i \mid \omega = j)}{\sum_{k \in \Omega} P(\omega = k) \prod_{i=1}^n P(s_i \mid \omega = k)} \\
&= \frac{\frac{1}{P(\omega=j)^{n-1}} \prod_{i=1}^n P(\omega = j) P(s_i \mid \omega = j)}{\sum_{k \in \Omega} \frac{1}{P(\omega=k)^{n-1}} \prod_{i=1}^n P(\omega = k) P(s_i \mid \omega = k)} \\
&= \frac{\frac{1}{P(\omega=j)^{n-1}} \prod_{i=1}^n \frac{P(\omega=j) P(s_i \mid \omega=j)}{P(s_i)}}{\sum_{k \in \Omega} \frac{1}{P(\omega=k)^{n-1}} \prod_{i=1}^n \frac{P(\omega=k) P(s_i \mid \omega=k)}{P(s_i)}} \\
&= \frac{\frac{1}{P(\omega=j)^{n-1}} \prod_{i=1}^n r_{ij}}{\sum_{k \in \Omega} \frac{1}{P(\omega=k)^{n-1}} \prod_{i=1}^n r_{ik}}.
\end{aligned}$$

Since the above expression depends only on r_i 's but not on s_i 's, we have:

$$f_j^*(\mathbf{r}) = P(\omega = j \mid \mathbf{r}) = P(\omega = j \mid \mathbf{s}) = \frac{\frac{1}{P(\omega=j)^{n-1}} \prod_{i=1}^n r_{ij}}{\sum_{k \in \Omega} \frac{1}{P(\omega=k)^{n-1}} \prod_{i=1}^n r_{ik}}.$$

(For the special binary event case,

$$f^*(\mathbf{r}) = P(\omega = 1 \mid \mathbf{r}) = \frac{\frac{1}{P(\omega=1)^{n-1}} \prod_{i=1}^n r_i}{\frac{1}{P(\omega=1)^{n-1}} \prod_{i=1}^n r_i + \frac{1}{P(\omega=0)^{n-1}} \prod_{i=1}^n (1 - r_i)} = \frac{1}{1 + \left(\frac{P(\omega=1)}{P(\omega=0)}\right)^{n-1} \prod_{i=1}^n \frac{1-r_i}{r_i}}.$$

This proves Lemma 5.1.) \square

According to Lemma C.2, the optimal aggregator $\mathbf{f}^* = (f_j)_{j \in \Omega}$ belongs to the following family of aggregators, parameterized by $\boldsymbol{\theta} = (\theta_j)_{j \in \Omega} \in \mathbb{R}_+^{|\Omega|}$:

$$\mathcal{F} = \left\{ \mathbf{f}^\theta = (f_j^\theta)_{j \in \Omega} \mid f_j^\theta(\mathbf{r}) = \frac{\theta_j \prod_{i=1}^n r_{ij}}{\sum_{k \in \Omega} \theta_k \prod_{i=1}^n r_{ik}} \right\} \quad (20)$$

(The optimal aggregator \mathbf{f}^* has parameter $\theta_j = \frac{1}{P(\omega=j)^{n-1}}$.) Let \mathcal{G} be the family of loss functions associated with \mathcal{F} , which is also parameterized by $\boldsymbol{\theta} \in \mathbb{R}_+^{|\Omega|}$,

$$\mathcal{G} = \left\{ g^\theta \mid g^\theta(\mathbf{r}, \omega) = \sum_{j \in \Omega} |f_j^\theta(\mathbf{r}) - \mathbb{1}[\omega = j]|^2 \right\}. \quad (21)$$

We recall the definition of pseudo-dimension of a family of functions:

Definition C.3 (e.g., [1]). *Let \mathcal{H} be a family of functions from input space \mathcal{X} to real numbers \mathbb{R} . Let $x^{(1)}, \dots, x^{(d)} \in \mathcal{X}$ be d inputs. Let $t^{(1)}, \dots, t^{(d)} \in \mathbb{R}$ be d "thresholds". Let $\mathbf{b} = (b^{(1)}, \dots, b^{(d)}) \in \{-1, +1\}^d$ be a vector of labels. We say \mathbf{b} can be generated by \mathcal{H} (on inputs $x^{(1)}, \dots, x^{(d)}$ with thresholds $t^{(1)}, \dots, t^{(d)}$) if there exists a function $h_{\mathbf{b}} \in \mathcal{H}$ such that $h_{\mathbf{b}}(x^{(i)}) > t^{(i)}$ if $b^{(i)} = 1$ and $h_{\mathbf{b}}(x^{(i)}) < t^{(i)}$ if $b^{(i)} = -1$ (namely, $\text{sign}(h_{\mathbf{b}}(x^{(i)}) - t^{(i)}) = b^{(i)}$) for every $i \in \{1, \dots, d\}$. The pseudo-dimension of \mathcal{H} , $\text{Pdim}(\mathcal{H})$, is the largest integer d for which there exist d inputs and d thresholds such that all the 2^d label vectors in $\{-1, +1\}^d$ can be generated by \mathcal{H} .*

Pseudo-dimension gives a sample complexity upper bound for the uniform convergence of the empirical means of a family of functions to their true means:

Theorem C.4 (e.g., [1]). *Let \mathcal{H} be a family of functions from \mathcal{X} to $[0, U]$ with $\text{Pdim}(\mathcal{H}) = d$. For any distribution \mathcal{D} on \mathcal{X} , with probability at least $1 - \delta$ over the random draw of*

$$T = O\left(\frac{U^2}{\varepsilon^2} \left(d \cdot \log \frac{U}{\varepsilon} + \log \frac{1}{\delta}\right)\right)$$

samples $x^{(1)}, \dots, x^{(T)}$ from \mathcal{D} , we have: for every $h \in \mathcal{H}$, $|\mathbb{E}_{x \sim \mathcal{D}}[h(x)] - \frac{1}{T} \sum_{t=1}^T h(x^{(t)})| \leq \varepsilon$.

We now upper bound the pseudo-dimension of \mathcal{G} :

Lemma C.5. $\text{Pdim}(\mathcal{G}) \leq O(|\Omega| \log |\Omega|)$.

Proof. Suppose $\text{Pdim}(\mathcal{G}) = d$. By definition, there exist d inputs $(\mathbf{r}^{(1)}, \omega^{(1)}), \dots, (\mathbf{r}^{(d)}, \omega^{(d)})$ and d thresholds $t^{(1)}, \dots, t^{(d)} \in \mathbb{R}$ such that all the 2^d label vectors $\mathbf{b} \in \{-1, +1\}^d$ can be generated by some function $g^{\theta_{\mathbf{b}}} \in \mathcal{G}$. We will count how many label vectors can actually be generated by \mathcal{G} .

Consider the ℓ -th input $(\mathbf{r}^{(\ell)}, \omega^{(\ell)})$ and threshold $t^{(\ell)}$. To simplify notations we omit superscript (ℓ) , so we have (\mathbf{r}, ω) and t . We write $\mathbf{x} = (x_j)_{j \in \Omega}$ where $x_j = \prod_{i=1}^n r_{ij}$. For any function $g^\theta \in \mathcal{G}$, we have

$$\begin{aligned} g^\theta(\mathbf{r}, \omega) &= \sum_{j \in \Omega} \left| \frac{\theta_j x_j}{\sum_{k \in \Omega} \theta_k x_k} - \mathbb{1}[\omega = j] \right|^2 \\ &= \sum_{j \in \Omega} \left(\frac{\theta_j x_j}{\sum_{k \in \Omega} \theta_k x_k} \right)^2 - 2 \frac{\theta_\omega x_\omega}{\sum_{k \in \Omega} \theta_k x_k} + 1 \\ &= \frac{1}{\left(\sum_{k \in \Omega} \theta_k x_k \right)^2} \left[\sum_{j \in \Omega} (\theta_j x_j)^2 - 2 \theta_\omega x_\omega \sum_{k \in \Omega} \theta_k x_k + \left(\sum_{k \in \Omega} \theta_k x_k \right)^2 \right]. \end{aligned}$$

By definition, the set of parameters that give input (\mathbf{r}, ω) label “+1” is $\{\theta \in \mathbb{R}_+^{|\Omega|} \mid g^\theta(\mathbf{r}, \omega) > t\}$. By the above equation, this set is equal to the set

$$\left\{ \theta \in \mathbb{R}_+^{|\Omega|} \mid \sum_{j \in \Omega} (\theta_j x_j)^2 - 2 \theta_\omega x_\omega \sum_{k \in \Omega} \theta_k x_k + \left(\sum_{k \in \Omega} \theta_k x_k \right)^2 > t \left(\sum_{k \in \Omega} \theta_k x_k \right)^2 \right\}.$$

We note that the above set is the solution set to the quadratic form inequality $\theta^\top A \theta > 0$ for some matrix $A \in \mathbb{R}^{|\Omega| \times |\Omega|}$. Similarly, the set of parameters that give input (\mathbf{r}, ω) label “-1” is the solution set to $\theta^\top A \theta < 0$. These two sets share a boundary: the solution set to $\theta^\top A \theta = 0$, which is a hyper-ellipsoid in $\mathbb{R}^{|\Omega|}$. In other words, the input (\mathbf{r}, ω) and threshold t define a hyper-ellipsoid which divides the parameter space $\mathbb{R}_+^{|\Omega|}$ into two regions such that all the parameters in one region generate the same label for that input.

Enumerating all inputs $(\mathbf{r}^{(1)}, \omega^{(1)}), \dots, (\mathbf{r}^{(d)}, \omega^{(d)})$. They define d hyper-ellipsoids in total, dividing the parameter space $\mathbb{R}_+^{|\Omega|}$ into several regions. Within each region, all the parameters generate the same label for each input and hence generate the same label vector. So, the number of label vectors that can be generated by all the parameters in $\mathbb{R}_+^{|\Omega|}$ is upper bounded by the number of regions. The number of regions divided by d hyper-ellipsoids in $\mathbb{R}_+^{|\Omega|}$ is in the order of $O(d^{|\Omega|})$. Hence, to generate all the 2^d label vectors we need $O(d^{|\Omega|}) \geq 2^d$. This gives $d \leq O(|\Omega| \log |\Omega|)$. \square

By Theorem C.4 and Lemma C.5, plugging in $d = \text{Pdim}(\mathcal{G}) \leq O(|\Omega| \log |\Omega|)$ and $U = 2$ (because $g^\theta(\mathbf{r}, \omega)$ is bounded by $[0, 2]$ according to (18)), we obtain the following: with probability at least $1 - \delta$ over the random draw of

$$T = O\left(\frac{U^2}{\varepsilon^2} \left(d \cdot \log \frac{U}{\varepsilon} + \log \frac{1}{\delta}\right)\right) = O\left(\frac{|\Omega| \log |\Omega|}{\varepsilon^2} \log \frac{1}{\varepsilon} + \frac{1}{\varepsilon^2} \log \frac{1}{\delta}\right)$$

samples $(\mathbf{r}^{(1)}, \omega^{(1)}), \dots, (\mathbf{r}^{(T)}, \omega^{(T)})$ from P , we have for any $g^\theta \in \mathcal{G}$,

$$\left| \mathbb{E}_P[g^\theta(\mathbf{r}, \omega)] - \frac{1}{T} \sum_{t=1}^T g^\theta(\mathbf{r}^{(t)}, \omega^{(t)}) \right| \leq \varepsilon.$$

This is equivalent to for any $\mathbf{f}^\theta \in \mathcal{F}$,

$$\begin{aligned} \left| \mathbb{E}_P \left[\sum_{j \in \Omega} |f_j^\theta(\mathbf{r}) - \mathbb{1}[\omega = j]|^2 \right] - \frac{1}{T} \sum_{t=1}^T \sum_{j \in \Omega} |f_j^\theta(\mathbf{r}^{(t)}) - \mathbb{1}[\omega^{(t)} = j]|^2 \right| &\leq \varepsilon \\ \iff \left| L_P(\mathbf{f}^\theta) - L_{\hat{P}_{\text{emp}}}(\mathbf{f}^\theta) \right| &\leq \varepsilon, \end{aligned}$$

where \hat{P}_{emp} is the empirical distribution $\text{Uniform}\{(\mathbf{r}^{(1)}, \omega^{(1)}), \dots, (\mathbf{r}^{(T)}, \omega^{(T)})\}$. By the same logic as the proof of the upper bound in Theorem 4.1 (Section 4.1), the empirically optimal aggregator $\hat{\mathbf{f}}^* = \text{argmin}_{\mathbf{f}^\theta \in \mathcal{F}} L_{\hat{P}_{\text{emp}}}(\mathbf{f}^\theta)$ is 2ε -optimal.

D Missing Proofs from Section 2

D.1 Proof of Lemma 2.1

To prove the first item $f^*(\mathbf{r}) = P(\omega = 1 \mid \mathbf{r})$, we simply note that

$$f^* = \underset{f}{\operatorname{argmin}} \mathbb{E}_{\mathbf{r}} \left[\mathbb{E}[|f(\mathbf{r}) - \omega|^2 \mid \mathbf{r}] \right]$$

should minimize $\mathbb{E}[|f(\mathbf{r}) - \omega|^2 \mid \mathbf{r}]$ for almost every \mathbf{r} . This gives $f^*(\mathbf{r}) = \mathbb{E}[\omega \mid \mathbf{r}] = P(\omega = 1 \mid \mathbf{r})$.

To prove the second item $L_P(f) - L_P(f^*) = \mathbb{E}_P[|f(\mathbf{r}) - f^*(\mathbf{r})|^2]$, we note that, by the definition of $L_P(\cdot)$ and the fact that $f^*(\mathbf{r}) = \mathbb{E}[\omega \mid \mathbf{r}]$ proven above,

$$\begin{aligned} L_P(f) - L_P(f^*) &= \mathbb{E}[|f(\mathbf{r}) - \omega|^2] - \mathbb{E}[|f^*(\mathbf{r}) - \omega|^2] \\ &= \mathbb{E}[f(\mathbf{r})^2] - 2\mathbb{E}[f(\mathbf{r})\omega] - \mathbb{E}[f^*(\mathbf{r})^2] + 2\mathbb{E}[f^*(\mathbf{r})\omega] \\ &= \mathbb{E}[f(\mathbf{r})^2] - 2\mathbb{E}_{\mathbf{r}}[f(\mathbf{r})\mathbb{E}[\omega \mid \mathbf{r}]] - \mathbb{E}[f^*(\mathbf{r})^2] + 2\mathbb{E}_{\mathbf{r}}[f^*(\mathbf{r})\mathbb{E}[\omega \mid \mathbf{r}]] \\ &= \mathbb{E}[f(\mathbf{r})^2] - 2\mathbb{E}_{\mathbf{r}}[f(\mathbf{r})f^*(\mathbf{r})] - \mathbb{E}[f^*(\mathbf{r})^2] + 2\mathbb{E}_{\mathbf{r}}[f^*(\mathbf{r})^2] \\ &= \mathbb{E}[f(\mathbf{r})^2 - 2f(\mathbf{r})f^*(\mathbf{r}) + f^*(\mathbf{r})^2] \\ &= \mathbb{E}[|f(\mathbf{r}) - f^*(\mathbf{r})|^2]. \end{aligned}$$

E Missing Proofs from Section 4

E.1 Proof of Lemma 4.3

Recall that we have a family \mathcal{D} of distributions over $\mathcal{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ satisfying the three properties in Definition 4.2 (B -uniformly bounded, same marginal across distributions, and distinct marginals across signals). We have constructed from \mathcal{D} the family of distributions $\mathcal{P} = \{P_D : D \in \mathcal{D}\}$ for the forecast aggregation problem as follows:

$$\begin{cases} P_D(\omega = 0) = P_D(\omega = 1) = \frac{1}{2}, \\ P_D(\cdot \mid \omega = 0) = \text{Uniform}(\mathcal{S}), & \text{namely, } P_D(\mathbf{s} \mid \omega = 0) = \frac{1}{|\mathcal{S}|} = \frac{1}{m^n}, \\ P_D(\cdot \mid \omega = 1) = D, & \text{namely, } P_D(\mathbf{s} \mid \omega = 1) = D(\mathbf{s}). \end{cases} \quad (22)$$

We want to show that ε -optimal aggregation with respect to \mathcal{P} will imply $(1+B)^2\sqrt{\varepsilon}$ total variation distance learning with respect to \mathcal{D} , and hence $T_{\mathcal{P}}(\varepsilon, \delta) \geq T_{\mathcal{D}}^{\text{TV}}((1+B)^2\sqrt{\varepsilon}, \delta)$.

First, we have the following observations about \mathcal{D} and \mathcal{P} :

Fact E.1. For a family of distributions \mathcal{D} that satisfies the three properties in Definition 4.2:

1. Given signal s_i , expert i 's report is $r_i = P_D(\omega = 1 \mid s_i) = \frac{D(s_i)}{\frac{1}{m} + D(s_i)}$, which is the same for all distributions $D \in \mathcal{D}$.
2. For every expert i , given different signals $s_i \neq s'_i$, its reports $r_i \neq r'_i$. So, there is a one-to-one mapping between s_i and r_i for every $i \in \{1, \dots, n\}$ and also a one-to-one mapping between the joint signal $\mathbf{s} = (s_1, \dots, s_n)$ and the joint report $\mathbf{r} = (r_1, \dots, r_n)$.
3. For any joint signal \mathbf{s} , with corresponding joint report \mathbf{r} , we have $D(\mathbf{s}) = \frac{f^*(\mathbf{r})}{m^n(1-f^*(\mathbf{r}))}$.

Proof. We prove the three items one by one:

1. By definition, the marginal distribution of joint signal, $P_D(\mathbf{s})$, is

$$P_D(\mathbf{s}) = P_D(\omega = 0)P_D(\mathbf{s} \mid \omega = 0) + P_D(\omega = 1)P_D(\mathbf{s} \mid \omega = 1) = \frac{1}{2} \left(\frac{1}{m^n} + D(\mathbf{s}) \right). \quad (23)$$

Fixing s_i , summing over $\mathbf{s}_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n) \in \mathcal{S}_{-i}$, we get

$$P_D(s_i) = \frac{1}{2} \left(\frac{|\mathcal{S}_{-i}|}{m^n} + \sum_{\mathbf{s}_{-i} \in \mathcal{S}_{-i}} D(\mathbf{s}) \right) = \frac{1}{2} \left(\frac{1}{m} + D(s_i) \right).$$

So, given signal s_i , expert i reports

$$r_i = P_D(\omega = 1 \mid s_i) = \frac{P_D(\omega = 1)P_D(s_i \mid \omega = 1)}{P_D(s_i)} = \frac{D(s_i)}{\frac{1}{m} + D(s_i)}, \quad (24)$$

and this is the same for all $D \in \mathcal{D}$ since $D(s_i) = D'(s_i)$ by the ‘‘same marginal across distributions’’ property.

2. Given $s_i \neq s'_i$, by the ‘‘distinct marginals across signals’’ property, we have $D(s_i) \neq D(s'_i)$. Since $r_i = \frac{D(s_i)}{\frac{1}{m} + D(s_i)}$ and $\frac{x}{\frac{1}{m} + x}$ is a strictly increasing function of x , it follows that $r_i \neq r'_i$.
3. According to item 2, there is a one-to-one mapping between $\mathbf{s} = (s_1, \dots, s_n)$ and $\mathbf{r} = (r_1, \dots, r_n)$; in other words, observing signals s_1, \dots, s_n is equivalent to observing reports r_1, \dots, r_n . Therefore, by Bayes’ rule we have

$$\begin{aligned} f^*(\mathbf{r}) &= P_D(\omega = 1 \mid \mathbf{r}) = P_D(\omega = 1 \mid \mathbf{s}) = \frac{P_D(\omega = 1)P_D(\mathbf{s} \mid \omega = 1)}{P_D(\mathbf{s})} \\ &\text{by (22) and (23)} = \frac{D(\mathbf{s})}{\frac{1}{m^n} + D(\mathbf{s})}. \end{aligned} \quad (25)$$

Rearranging, we obtain $D(\mathbf{s}) = \frac{f^*(\mathbf{r})}{m^n(1-f^*(\mathbf{r}))}$. \square

Claim E.2. *If we have an aggregator \hat{f} that is ε -optimal with respect to P_D , then we can find a distribution \hat{D} such that $d_{\text{TV}}(\hat{D}, D) \leq (1+B)^2\sqrt{\varepsilon}$.*

Proof. Because D is B -uniformly bounded, from (25) we can verify that $f^*(\mathbf{r})$ satisfies $f^*(\mathbf{r}) \leq \frac{B}{1+B}$. So, we can assume $\hat{f}(\mathbf{r}) \leq \frac{B}{1+B}$ as well (if $\hat{f}(\mathbf{r}) > \frac{B}{1+B}$, we can let $\hat{f}(\mathbf{r})$ be $\frac{B}{1+B}$; this only reduces the approximation error $\mathbb{E}[|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|^2]$). Define \hat{D} by letting $\hat{D}(\mathbf{s}) = \frac{\hat{f}(\mathbf{r})}{m^n(1-\hat{f}(\mathbf{r}))}$, $\forall \mathbf{s} \in \mathcal{S}$, where \mathbf{r} is the reports corresponding to \mathbf{s} (cf., Fact E.1). Then, $d_{\text{TV}}(\hat{D}, D)$ is

$$\begin{aligned} d_{\text{TV}}(\hat{D}, D) &= \frac{1}{2} \sum_{\mathbf{s} \in \mathcal{S}} |\hat{D}(\mathbf{s}) - D(\mathbf{s})| = \frac{1}{2} \sum_{\mathbf{s} \in \mathcal{S}} \left| \frac{\hat{f}(\mathbf{r})}{m^n(1-\hat{f}(\mathbf{r}))} - \frac{f^*(\mathbf{r})}{m^n(1-f^*(\mathbf{r}))} \right| \\ &= \frac{1}{2m^n} \sum_{\mathbf{s} \in \mathcal{S}} \left| \frac{\hat{f}(\mathbf{r})}{1-\hat{f}(\mathbf{r})} - \frac{f^*(\mathbf{r})}{1-f^*(\mathbf{r})} \right|. \end{aligned}$$

Because $\hat{f}(\mathbf{r}), f^*(\mathbf{r}) \leq \frac{B}{1+B}$ and the function $\frac{x}{1-x}$ has derivative $\frac{1}{(1-x)^2} \leq (1+B)^2$ when $x \leq \frac{B}{1+B}$, we have

$$\begin{aligned} d_{\text{TV}}(\hat{D}, D) &\leq \frac{1}{2m^n} (1+B)^2 \sum_{\mathbf{s} \in \mathcal{S}} |\hat{f}(\mathbf{r}) - f^*(\mathbf{r})| \\ &\text{by (23)} \leq (1+B)^2 \sum_{\mathbf{s} \in \mathcal{S}} P_D(\mathbf{s}) |\hat{f}(\mathbf{r}) - f^*(\mathbf{r})| = (1+B)^2 \mathbb{E}_{P_D} [|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|]. \end{aligned}$$

By Jensen’s inequality $(\mathbb{E}[X])^2 \leq \mathbb{E}[X^2]$ and by the assumption that \hat{f} is ε -optimal, we have $(\mathbb{E}_{P_D} [|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|])^2 \leq \mathbb{E}_{P_D} [|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|^2] \leq \varepsilon$. Thus,

$$d_{\text{TV}}(\hat{D}, D) \leq (1+B)^2\sqrt{\varepsilon},$$

which proves the claim. \square

Now, we present the reduction from learning \mathcal{D} in total variation distance to forecast aggregation for \mathcal{P} . We use notations $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(T)} \in \mathcal{S}$ to represent the samples from D . From $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(T)}$ we construct the samples $(\mathbf{r}^{(1)}, \omega^{(1)}), \dots, (\mathbf{r}^{(T)}, \omega^{(T)})$ for the forecast aggregation problem. After obtaining a solution \hat{f} to the latter problem, we convert it into a solution \hat{D} to the former. Details are as follows:

Input: T i.i.d. samples $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(T)}$ from an unknown distribution $D \in \mathcal{D}$.

Reduction:

1. Draw T samples $\omega^{(1)}, \dots, \omega^{(T)} \sim \text{Uniform}\{0, 1\}$.
2. For each $t = 1, \dots, T$, do the following:
 - If $\omega^{(t)} = 0$, draw $\mathbf{s}^{(t)} \sim \text{Uniform}(\mathcal{S})$.
 - If $\omega^{(t)} = 1$, let $\mathbf{s}^{(t)} = \mathbf{x}^{(t)}$.
 - For each i , compute $r_i^{(t)} = \frac{D(s_i^{(t)})}{\frac{1}{m} + D(s_i^{(t)})}$. Let $\mathbf{r}^{(t)} = (r_1^{(t)}, \dots, r_n^{(t)})$.
3. Feed $S_T = \{(\mathbf{r}^{(1)}, \omega^{(1)}), \dots, (\mathbf{r}^{(T)}, \omega^{(T)})\}$ to the forecast aggregation problem. Obtain solution \hat{f} .
4. Convert \hat{f} into \hat{D} according to Claim E.2.

Output: \hat{D} .

We remark that, in the second step of the reduction, the report $r_i^{(t)}$ can be computed even if D is unknown, because the $D(s_i^{(t)})$ is the same for all $D \in \mathcal{D}$ and hence known.

Using the above reduction, we show that the sample complexity of ε -optimal forecast aggregation for \mathcal{P} cannot be smaller than the sample complexity of learning \mathcal{D} within total variation distance $(1 + B)^2 \sqrt{\varepsilon}$, which will prove Lemma 4.3:

Proof of Lemma 4.3: First, we verify that the distribution of samples S_T in the above reduction is exactly the distribution of T samples $\{(\mathbf{r}^{(1)}, \omega^{(1)}), \dots, (\mathbf{r}^{(T)}, \omega^{(T)})\}$ from P_D . This is because: (1) the distribution of $\omega^{(t)}$ is $\text{Uniform}\{0, 1\}$, as defined in P_D ; (2) given $\omega^{(t)} = 0$, the distribution of $\mathbf{s}^{(t)}$ is $\text{Uniform}(\mathcal{S})$, as defined in P_D ; (3) given $\omega^{(t)} = 1$, the distribution of $\mathbf{s}^{(t)}$ is the same as the distribution of $\mathbf{x}^{(t)}$, which is D , because the random draws of $\omega^{(t)}$ and $\mathbf{x}^{(t)}$ are independent; (4) according to Fact E.1, the report $r_i^{(t)} = \frac{D(s_i^{(t)})}{\frac{1}{m} + D(s_i^{(t)})} = P_D(\omega = 1 \mid s_i^{(t)})$, as desired.

Then, by the definition of sample complexity of forecast aggregation, if we are given $T = T_{\mathcal{P}}(\varepsilon, \delta)$ samples S_T for the forecast aggregation problem, then with probability at least $1 - \delta$ we should be able to find an ε -optimal aggregator \hat{f} with respect to P_D . According to Claim E.2, we can convert \hat{f} into a \hat{D} such that

$$d_{\text{TV}}(\hat{D}, D) \leq (1 + B)^2 \sqrt{\varepsilon}.$$

By the definition of sample complexity $T_{\mathcal{D}}^{\text{TV}}(\cdot, \delta)$ of distribution learning, T must be at least

$$T \geq T_{\mathcal{D}}^{\text{TV}}((1 + B)^2 \sqrt{\varepsilon}, \delta),$$

which proves the lemma.

E.2 Proof of Proposition 4.4

To prove Proposition 4.4 we will construct a family of distributions \mathcal{D} that satisfies the three properties in Definition 4.2 and requires $T_{\mathcal{D}}^{\text{TV}}(\varepsilon_{\text{TV}}, \delta) = \Omega\left(\frac{m^{n-2} + \log(1/\delta)}{\varepsilon_{\text{TV}}^2}\right)$ samples to learn. For simplicity, we write $\varepsilon = \varepsilon_{\text{TV}}$. For technical convenience, we assume $\varepsilon < \frac{1}{40}$, $\delta < 0.01$.

E.2.1 Part 1: Constructing \mathcal{D}

We index the distributions $D_z \in \mathcal{D}$ by z ; the meaning of z will be defined later. Without loss of generality, we assume $|\mathcal{S}_i| = m$ to be an even integer, and denote $\mathcal{S}_i = \{1, \dots, m\} =: S$. We will define D_z to be a distribution over the joint signal space $\mathcal{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_n = S^n = \{1, \dots, m\}^n$. In the following we will call a joint signal $\mathbf{s} \in S^n$ simply a signal. We write a signal $\mathbf{s} \in S^n$ as $\mathbf{s} = (\mathbf{b}, x, y)$, where $\mathbf{b} \in S^{n-2}$ and $x, y \in S$. We sort all the m^n signals in S^n by the lexicographical

order, from $(1, \dots, 1, 1)$, $(1, \dots, 1, 2)$, ..., to (m, \dots, m, m) . We number the signals from 1 to m^n , using

$$\text{num}(\mathbf{s}) = \text{num}(\mathbf{b}, x, y) \in \{1, \dots, m^n\}.$$

to denote their numbers. We divide the whole signal space S^n into $|S^{n-2}| = m^{n-2}$ ‘‘buckets’’, each of size m^2 and denoted by $B_{\mathbf{b}}$:

$$B_{\mathbf{b}} = \{(\mathbf{b}, x, y) : x, y \in S\}, \quad \mathbf{b} \in S^{n-2}.$$

We first define a ‘‘base’’ distribution D_{base} , then construct the distributions D_z 's by modifying the probabilities of the base distribution within each bucket $B_{\mathbf{b}}$. Let $\gamma = 1 + \frac{1}{m^n}$. The base distribution is defined as follows:

$$D_{\text{base}}(\mathbf{s}) = \frac{\gamma^{\text{num}(\mathbf{s})}}{W}, \quad W = \sum_{\mathbf{s} \in S^n} \gamma^{\text{num}(\mathbf{s})} = \sum_{\ell=1}^{m^n} \gamma^\ell.$$

Because $1 \leq \gamma^\ell \leq \gamma^{m^n} \leq (1 + \frac{1}{m^n})^{m^n} \leq e$, we have

$$m^n \leq W \leq em^n. \quad (26)$$

We assign a sign $z_{\mathbf{b}} \in \{+1, -1\}$ to each bucket $B_{\mathbf{b}}$, and let z be a vector of length m^{n-2} that includes the signs of all buckets:

$$z = (z_{\mathbf{b}})_{\mathbf{b} \in S^{n-2}}, \quad z_{\mathbf{b}} \in \{+1, -1\}.$$

We have $2^{m^{n-2}}$ different z 's in total, and hence $2^{m^{n-2}}$ different distributions D_z 's in \mathcal{D} in total. Let $c = 20$, so that $c\varepsilon \leq 1/2$. For each z , we define D_z as follows: within each bucket $B_{\mathbf{b}}$, for each element $(\mathbf{b}, x, y) \in B_{\mathbf{b}}$ let

$$D_z(\mathbf{b}, x, y) = \begin{cases} D_{\text{base}}(\mathbf{b}, x, y) + \frac{z_{\mathbf{b}} c \varepsilon}{W} & \text{if } x \leq \frac{m}{2}, y \leq \frac{m}{2} \\ D_{\text{base}}(\mathbf{b}, x, y) - \frac{z_{\mathbf{b}} c \varepsilon}{W} & \text{if } x \leq \frac{m}{2}, y > \frac{m}{2} \\ D_{\text{base}}(\mathbf{b}, x, y) - \frac{z_{\mathbf{b}} c \varepsilon}{W} & \text{if } x > \frac{m}{2}, y \leq \frac{m}{2} \\ D_{\text{base}}(\mathbf{b}, x, y) + \frac{z_{\mathbf{b}} c \varepsilon}{W} & \text{if } x > \frac{m}{2}, y > \frac{m}{2} \end{cases}. \quad (27)$$

Claim E.3. *The family of distributions $\mathcal{D} = \{D_z\}_z$ defined above satisfies the three properties in Definition 4.2: B -uniformly bounded with $B = e + 1/2$, same marginal across distributions, distinct marginals across signals.*

Proof. **B -uniformly bounded:** For any \mathbf{s} , any D_z , by definition,

$$D_z(\mathbf{s}) \leq D_{\text{base}}(\mathbf{s}) + \frac{c\varepsilon}{W} \leq \frac{\gamma^{m^n}}{W} + \frac{c\varepsilon}{W} \stackrel{(26)}{\leq} \frac{e}{m^n} + \frac{1/2}{m^n}.$$

So, the distribution is B -uniformly bounded with $B = e + 1/2$.

Same marginal across distributions: Consider each $D_z(s_i)$. We want to show that $D_z(s_i)$ does not depend on z , and in fact, $D_z(s_i) = D_{\text{base}}(s_i)$. If $i \in \{1, \dots, n-2\}$, namely, s_i is a component of the vector \mathbf{b} , then we have

$$D_z(s_i) = \sum_{\mathbf{s}_{-i} \in \mathcal{S}_{-i}} D_z(s_i, \mathbf{s}_{-i}) = \sum_{\mathbf{b} \in S^{n-2}: b_i = s_i} \sum_{x=1}^m \sum_{y=1}^m D_z(\mathbf{b}, x, y).$$

We notice that, fixing any \mathbf{b} , the numbers of $\frac{+z_{\mathbf{b}} c \varepsilon}{W}$ and $\frac{-z_{\mathbf{b}} c \varepsilon}{W}$ in the summation $\sum_{x=1}^m \sum_{y=1}^m D_z(\mathbf{b}, x, y)$ are the same. So, they cancel out, and we obtain

$$D_z(s_i) = \sum_{\mathbf{b} \in S^{n-2}: b_i = s_i} \sum_{x=1}^m \sum_{y=1}^m D_{\text{base}}(\mathbf{b}, x, y) = D_{\text{base}}(s_i).$$

If $i = n-1$, namely $s_i = x$, then we have:

$$D_z(s_i) = \sum_{\mathbf{s}_{-i} \in \mathcal{S}_{-i}} D_z(s_i, \mathbf{s}_{-i}) = \sum_{\mathbf{b} \in S^{n-2}} \sum_{y=1}^m D_z(\mathbf{b}, x, y).$$

Fixing any \mathbf{b} , the numbers of $\frac{+z_{\mathbf{b}}c\varepsilon}{W}$ and $\frac{-z_{\mathbf{b}}c\varepsilon}{W}$ in the summation $\sum_{y=1}^m D_z(\mathbf{b}, x, y)$ are the same. So, they cancel out, and we obtain

$$D_z(s_i) = \sum_{\mathbf{b} \in S^{n-2}} \sum_{y=1}^m D_{\text{base}}(\mathbf{b}, x, y) = D_{\text{base}}(s_i).$$

Finally, if $i = n$, namely $s_i = y$, then by a similar argument as above we have

$$D_z(s_i) = \sum_{\mathbf{b} \in S^{n-2}} \sum_{x=1}^m D_{\text{base}}(\mathbf{b}, x, y) = D_{\text{base}}(s_i).$$

Distinct marginals across signals: By the ‘‘same marginal across distributions’’ property above we have $D_z(s_i) = D_{\text{base}}(s_i)$. So, to prove ‘‘distinct marginals across signals’’ we only need to prove $D_{\text{base}}(s_i) \neq D_{\text{base}}(s'_i)$ for $s_i \neq s'_i$. Without loss of generality assume $s_i < s'_i$. By the definition $D_{\text{base}}(s) = \frac{\gamma^{\text{num}(s)}}{W}$ and the fact that $\text{num}(s_i, \mathbf{s}_{-i}) < \text{num}(s'_i, \mathbf{s}_{-i})$ for any $\mathbf{s}_{-i} \in \mathcal{S}_{-i}$, we have

$$D_{\text{base}}(s_i) = \sum_{\mathbf{s}_{-i} \in \mathcal{S}_{-i}} D_{\text{base}}(s_i, \mathbf{s}_{-i}) < \sum_{\mathbf{s}_{-i} \in \mathcal{S}_{-i}} D_{\text{base}}(s'_i, \mathbf{s}_{-i}) = D_{\text{base}}(s'_i),$$

so $D_{\text{base}}(s_i) \neq D_{\text{base}}(s'_i)$. \square

E.2.2 Part 2: Sample Complexity Lower Bound of Learning \mathcal{D}

Overview We want to prove the proposition that the sample complexity of learning the family of distributions $D = \{D_z\}_z$ defined above is at least $T_{\mathcal{D}}^{\text{TV}}(\varepsilon, \delta) = \Omega\left(\frac{m^{n-2} + \log(1/\delta)}{\varepsilon^2}\right)$. This proof is analogous to a textbook proof of Proposition 3.3 (the sample complexity for learning *all* distributions), which uses reductions from the *distinguishing distributions* problem. Roughly speaking, if one can learn the unknown distribution D_z well then one must be able to guess most of the components of the sign vector $z = (z_{\mathbf{b}})_{\mathbf{b} \in S^{n-2}}$ correctly, meaning that one can distinguish whether the distribution $D_{z_{\mathbf{b}}}$ on bucket $B_{\mathbf{b}}$ is $D_{z_{\mathbf{b}}=+1}$ or $D_{z_{\mathbf{b}}=-1}$. However, since $D_{z_{\mathbf{b}}=+1}$ and $D_{z_{\mathbf{b}}=-1}$ are ‘‘ $O(\varepsilon)$ -close’’ to each other, distinguishing them requires $\Omega(\frac{1}{\varepsilon^2})$ samples. In average, there are only $O(\frac{T}{m^{n-2}})$ samples falling into a bucket (because there are m^{n-2} buckets in total and the distribution D_z is close to uniform). We thus need $O(\frac{T}{m^{n-2}}) = \Omega(\frac{1}{\varepsilon^2})$, which gives $T = \Omega(\frac{m^{n-2}}{\varepsilon^2})$. Ignoring logarithmic terms, this proves the proposition.

Formal argument First, we note that if we can learn D_z very well, then we can guess the vector z correctly for a large fraction of its m^{n-2} components. Formally, suppose we obtain from T samples a distribution \hat{D} such that $d_{\text{TV}}(\hat{D}, D_z) \leq \varepsilon$. We find the distribution D_w in \mathcal{D} , $w = (w_{\mathbf{b}})_{\mathbf{b} \in S^{n-2}} \in \{+1, -1\}^{m^{n-2}}$, that is closest to \hat{D} in total variation distance. By definition, we have $d_{\text{TV}}(D_w, \hat{D}) \leq d_{\text{TV}}(D_z, \hat{D}) \leq \varepsilon$. Hence, by triangle inequality,

$$d_{\text{TV}}(D_w, D_z) \leq d_{\text{TV}}(D_w, \hat{D}) + d_{\text{TV}}(\hat{D}, D_z) \leq 2\varepsilon.$$

Let

$$\# = |\{\mathbf{b} \in S^{n-2} \mid w_{\mathbf{b}} \neq z_{\mathbf{b}}\}|$$

be the number of different components of w and z . We claim that:

Claim E.4. $d_{\text{TV}}(D_w, D_z) \leq 2\varepsilon$ implies $\# \leq \frac{2W}{cm^2}$.

Proof. Whenever we have a different component $w_{\mathbf{b}} \neq z_{\mathbf{b}}$, this different component contributes the following to the total variation distance between D_w and D_z :

$$\frac{1}{2} \sum_{(\mathbf{b}, x, y) \in B_{\mathbf{b}}} |D_w(\mathbf{b}, x, y) - D_z(\mathbf{b}, x, y)| = \frac{1}{2} \sum_{(\mathbf{b}, x, y) \in B_{\mathbf{b}}} \frac{2c\varepsilon}{W} = \frac{m^2 c\varepsilon}{W}. \quad (28)$$

So, the number of different components of w and z is at most $\frac{d_{\text{TV}}(D_w, D_z)}{\frac{m^2 c\varepsilon}{W}} \leq \frac{2W}{cm^2}$. \square

We first show the $\Omega(\frac{\log(1/\delta)}{\varepsilon^2})$ part in the sample complexity lower bound, and then show the $\Omega(\frac{m^{n-2}}{\varepsilon^2})$ part. Together, they give the lower bound $\max\{\Omega(\frac{\log(1/\delta)}{\varepsilon^2}), \Omega(\frac{m^{n-2}}{\varepsilon^2})\} = \Omega(\frac{m^{n-2} + \log(1/\delta)}{\varepsilon^2})$. Consider the distribution $D_{+1} \in \mathcal{D}$ whose index is the all “+1” vector and the distribution $D_{-1} \in \mathcal{D}$ whose index is the all “-1” vector. According to (28), the total variation distance between D_{+1} and D_{-1} is $d_{\text{TV}}(D_{+1}, D_{-1}) = m^{n-2} \cdot \frac{m^2 c\varepsilon}{W} = \frac{m^n c\varepsilon}{W}$ because +1 and -1 have m^{n-2} different components. Since $W \leq em^n$ (26), we have $d_{\text{TV}}(D_{+1}, D_{-1}) \geq \frac{c\varepsilon}{e} > 2\varepsilon$. Consider the distinguishing distributions problem (defined in Section 3) where we want to distinguish D_{+1} and D_{-1} . If we can learn from samples a distribution \hat{D} that is ε -close in total variation distance to the unknown distribution D_{+1} or D_{-1} , then we can perfectly tell whether the unknown distribution is D_{+1} or D_{-1} because the two distributions are more than 2ε -away from each other in total variation distance. Lemma A.5 implies that, to distinguish D_{+1} and D_{-1} with probability $1 - \delta$, the number of samples must be at least $\Omega(\frac{\log(1/\delta)}{d_{\text{H}}^2(D_{+1}, D_{-1})}) = \Omega(\frac{\log(1/\delta)}{\varepsilon^2})$. This proves the $\Omega(\frac{\log(1/\delta)}{\varepsilon^2})$ part.

We then prove the $\Omega(\frac{m^{n-2}}{\varepsilon^2})$ part. Suppose we first draw the vector z from $\{+1, -1\}^{m^{n-2}}$ uniformly at random, then draw T samples from D_z . We obtain the D_w as above. Let’s consider the expected number of different components of w and z in this two-step random draw procedure:

$$\mathbb{E}_{z, T \text{ samples}} [\#\] = \mathbb{E} \left[\sum_{\mathbf{b} \in S^{n-2}} \mathbb{1}\{z_{\mathbf{b}} \neq w_{\mathbf{b}}\} \right] = \sum_{\mathbf{b} \in S^{n-2}} \mathbb{E} [\mathbb{1}\{z_{\mathbf{b}} \neq w_{\mathbf{b}}\}]. \quad (29)$$

We consider each component $\mathbb{E}[\mathbb{1}\{z_{\mathbf{b}} \neq w_{\mathbf{b}}\}]$ in the above summation. Suppose that, within the T samples drawn from D_z , $T_{\mathbf{b}}$ of them fall into the bucket $B_{\mathbf{b}}$. So, $T_{\mathbf{b}}$ follows the Binomial($T, D(B_{\mathbf{b}})$) distribution with

$$D(B_{\mathbf{b}}) = \sum_{(\mathbf{b}, x, y) \in B_{\mathbf{b}}} D_z(\mathbf{b}, x, y) = \frac{1}{W} \sum_{(\mathbf{b}, x, y) \in B_{\mathbf{b}}} \gamma^{\text{num}(\mathbf{b}, x, y)}.$$

(Notice that the $+\frac{z_{\mathbf{b}} c\varepsilon}{W}$ and $-\frac{z_{\mathbf{b}} c\varepsilon}{W}$ cancel out in the summation and hence $D(B_{\mathbf{b}})$ does not depend on z .) Let $D_{z_{\mathbf{b}}}$ denote the “ $B_{\mathbf{b}}$ -part” of distribution D_z , namely, D_z conditioning on $B_{\mathbf{b}}$:

$$D_{z_{\mathbf{b}}}(\mathbf{s}) = \frac{D_z(\mathbf{s})}{D(B_{\mathbf{b}})}, \quad \forall \mathbf{s} \in B_{\mathbf{b}}.$$

We think of the random draw of the vector z and the T samples as follows: first, we draw $T_{\mathbf{b}}$ from Binomial($T, D(B_{\mathbf{b}})$); second, we draw $z_{\mathbf{b}} \in \{+1, -1\}$ uniformly at random; third, we draw $T_{\mathbf{b}}$ samples from the conditional distribution $D_{z_{\mathbf{b}}}$; fourth, we draw the remaining vector $z_{-\mathbf{b}}$ and the remaining $T - T_{\mathbf{b}}$ samples (which are samples outside of $B_{\mathbf{b}}$). Only writing the first two steps explicitly, we have

$$\begin{aligned} \mathbb{E}[\mathbb{1}\{z_{\mathbf{b}} \neq w_{\mathbf{b}}\}] &= \mathbb{E}_{T_{\mathbf{b}}} \left[\mathbb{E}_{z_{\mathbf{b}}} \left[\mathbb{E}[\mathbb{1}\{z_{\mathbf{b}} \neq w_{\mathbf{b}}\} \mid z_{\mathbf{b}}, T_{\mathbf{b}}] \right] \right] \\ &= \mathbb{E}_{T_{\mathbf{b}}} \left[\frac{1}{2} \mathbb{E}[\mathbb{1}\{z_{\mathbf{b}} \neq w_{\mathbf{b}}\} \mid z_{\mathbf{b}} = +1, T_{\mathbf{b}}] + \frac{1}{2} \mathbb{E}[\mathbb{1}\{z_{\mathbf{b}} \neq w_{\mathbf{b}}\} \mid z_{\mathbf{b}} = -1, T_{\mathbf{b}}] \right] \\ &= \mathbb{E}_{T_{\mathbf{b}}} \left[\frac{1}{2} \Pr[z_{\mathbf{b}} \neq w_{\mathbf{b}} \mid z_{\mathbf{b}} = +1, T_{\mathbf{b}}] + \frac{1}{2} \Pr[z_{\mathbf{b}} \neq w_{\mathbf{b}} \mid z_{\mathbf{b}} = -1, T_{\mathbf{b}}] \right]. \quad (30) \end{aligned}$$

Claim E.5. For any $T_{\mathbf{b}}$, $\frac{1}{2} \Pr[z_{\mathbf{b}} \neq w_{\mathbf{b}} \mid z_{\mathbf{b}} = +1, T_{\mathbf{b}}] + \frac{1}{2} \Pr[z_{\mathbf{b}} \neq w_{\mathbf{b}} \mid z_{\mathbf{b}} = -1, T_{\mathbf{b}}] \geq \frac{1}{2} - 2c\varepsilon\sqrt{T_{\mathbf{b}}}$.

Proof. We notice that $\frac{1}{2} \Pr[z_{\mathbf{b}} \neq w_{\mathbf{b}} \mid z_{\mathbf{b}} = +1, T_{\mathbf{b}}] + \frac{1}{2} \Pr[z_{\mathbf{b}} \neq w_{\mathbf{b}} \mid z_{\mathbf{b}} = -1, T_{\mathbf{b}}]$ is the probability that we make a mistake when guessing the sign $z_{\mathbf{b}}$ using $w_{\mathbf{b}}$, if (1) $z_{\mathbf{b}}$ is chosen from $\{-1, +1\}$ uniformly at random; (2) we are given $T_{\mathbf{b}}$ samples from $D_{z_{\mathbf{b}}}$; (3) we then draw the remaining vector $z_{-\mathbf{b}}$ and the remaining samples; (4) finally, we use the D_w computed from all samples to get $w_{\mathbf{b}}$. The steps (3) and (4) define a randomized function that maps the $T_{\mathbf{b}}$ samples of $D_{z_{\mathbf{b}}}$ to $w_{\mathbf{b}} \in \{0, 1\}$, and therefore, according to the first item of Lemma A.5, we have

$$\frac{1}{2} \Pr[z_{\mathbf{b}} \neq w_{\mathbf{b}} \mid z_{\mathbf{b}} = +1, T_{\mathbf{b}}] + \frac{1}{2} \Pr[z_{\mathbf{b}} \neq w_{\mathbf{b}} \mid z_{\mathbf{b}} = -1, T_{\mathbf{b}}] \geq \frac{1}{2} - \sqrt{\frac{T_{\mathbf{b}}}{2}} d_{\text{H}}(D_{z_{\mathbf{b}}=+1}, D_{z_{\mathbf{b}}=-1}). \quad (31)$$

Then we consider $d_H(D_{z_b=+1}, D_{z_b=-1})$. We use Lemma A.2 to do so. For any $\mathbf{s} \in B_b$, we have, on the one hand,

$$\frac{D_{z_b=+1}(\mathbf{s})}{D_{z_b=-1}(\mathbf{s})} = \frac{\gamma^{\text{num}(\mathbf{s}) \pm c\varepsilon}}{\gamma^{\text{num}(\mathbf{s}) \mp c\varepsilon}} \geq \frac{\gamma^{\text{num}(\mathbf{s}) - c\varepsilon}}{\gamma^{\text{num}(\mathbf{s}) + c\varepsilon}} \geq 1 - 2c\varepsilon,$$

because $\frac{a-b}{a+b} = 1 - \frac{2b}{a+b} \geq 1 - 2b$ for $a = \gamma^{\text{num}(\mathbf{s})} \geq 1$. On the other hand,

$$\frac{D_{z_b=+1}(\mathbf{s})}{D_{z_b=-1}(\mathbf{s})} \leq \frac{\gamma^{\text{num}(\mathbf{s}) + c\varepsilon}}{\gamma^{\text{num}(\mathbf{s}) - c\varepsilon}} \leq 1 + 4c\varepsilon,$$

because $\frac{a+b}{a-b} = 1 + \frac{2b}{a-b} \leq 1 + 4b$ for $a = \gamma^{\text{num}(\mathbf{s})} \geq 1$ and $b = c\varepsilon \leq \frac{1}{2}$. Therefore, by Lemma A.2 we have

$$d_H^2(D_{z_b=+1}, D_{z_b=-1}) \leq \frac{1}{2}(4c\varepsilon)^2 = 8c^2\varepsilon^2. \quad (32)$$

Combining (31) and (32) proves our claim. \square

By (30) and Claim E.5, we have $\mathbb{E}[\mathbb{1}\{z_b \neq w_b\}] \geq \mathbb{E}\left[\frac{1}{2} - 2c\varepsilon\sqrt{T_b}\right]$. Summing over $\mathbf{b} \in S^{n-2}$, (29) becomes

$$\begin{aligned} \sum_{\mathbf{b} \in S^{n-2}} \mathbb{E}[\mathbb{1}\{z_b \neq w_b\}] &\geq \sum_{\mathbf{b} \in S^{n-2}} \mathbb{E}\left[\frac{1}{2} - 2c\varepsilon\sqrt{T_b}\right] = \frac{m^{n-2}}{2} - 2c\varepsilon \sum_{\mathbf{b} \in S^{n-2}} \mathbb{E}[\sqrt{T_b}] \\ &\text{(by Jensen's inequality } \mathbb{E}[\sqrt{X}] \leq \sqrt{\mathbb{E}[X]}) \geq \frac{m^{n-2}}{2} - 2c\varepsilon \sum_{\mathbf{b} \in S^{n-2}} \sqrt{\mathbb{E}[T_b]}. \end{aligned}$$

Because $\mathbb{E}[T_b] = T \cdot D(B_b) = \frac{T}{W} \sum_{(b,x,y) \in B_b} \gamma^{\text{num}(b,x,y)} \leq \frac{T}{W} m^2 \gamma^{m^n}$, we have

$$\sum_{\mathbf{b} \in S^{n-2}} \mathbb{E}[\mathbb{1}\{z_b \neq w_b\}] \geq \frac{m^{n-2}}{2} - 2c\varepsilon \sum_{\mathbf{b} \in S^{n-2}} \sqrt{\frac{T}{W} m^2 \gamma^{m^n}} = \frac{m^{n-2}}{2} - 2c\varepsilon m^{n-2} \sqrt{\frac{T}{W} m^2 \gamma^{m^n}}. \quad (33)$$

Now, let's consider the probability with which we can obtain \hat{D} such that $d_{\text{TV}}(\hat{D}, D_z) \leq \varepsilon$. We will show that this probability is at most $0.99 < 1 - \delta$ if T is less than $10^{-5} \cdot \frac{m^{n-2}}{\varepsilon^2}$. Recall from Claim E.4 that $d_{\text{TV}}(\hat{D}, D_z) \leq \varepsilon$ implies $\# = \sum_{\mathbf{b} \in S^{n-2}} \mathbb{1}\{z_b \neq w_b\} \leq \frac{2W}{cm^2}$. So, the probability is at most

$$\begin{aligned} \Pr \left[\sum_{\mathbf{b} \in S^{n-2}} \mathbb{1}\{z_b \neq w_b\} \leq \frac{2W}{cm^2} \right] &= \Pr \left[\sum_{\mathbf{b} \in S^{n-2}} \mathbb{1}\{z_b = w_b\} \geq m^{n-2} - \frac{2W}{cm^2} \right] \\ &\text{(by Markov's inequality)} \leq \frac{\mathbb{E}[\sum_{\mathbf{b} \in S^{n-2}} \mathbb{1}\{z_b = w_b\}]}{m^{n-2} - \frac{2W}{cm^2}} \\ &\text{(by (33))} \leq \frac{\frac{m^{n-2}}{2} + 2c\varepsilon m^{n-2} \sqrt{\frac{T}{W} m^2 \gamma^{m^n}}}{m^{n-2} - \frac{2W}{cm^2}} \\ &\text{(\(\gamma^{m^n} \leq e\) and \(m^n \leq W \leq em^n\) by (26))} \leq \frac{\frac{m^{n-2}}{2} + 2c\varepsilon m^{n-2} \sqrt{\frac{eT}{m^{n-2}}}}{m^{n-2} - \frac{2em^n}{cm^2}} \\ &= \frac{\frac{1}{2} + 2c\sqrt{\frac{e\varepsilon^2}{m^{n-2}}} T}{1 - \frac{2e}{c}} < 0.99 < 1 - \delta, \end{aligned}$$

when $c = 20$, $T < 10^{-5} \cdot \frac{m^{n-2}}{\varepsilon^2}$, and $\delta < 0.01$. This means that, in order to obtain such \hat{D} with probability at least $1 - \delta$, at least $10^{-5} \cdot \frac{m^{n-2}}{\varepsilon^2}$ samples are needed.

F Missing Proofs from Section 5

F.1 Proof of Theorem 5.2: the $\Omega(\frac{1}{\varepsilon} \log \frac{1}{\delta})$ Lower Bound

The proof uses a reduction from the distinguishing distributions problem (defined in Section 3). We construct two conditionally independent distributions P^1, P^2 over the space $\Omega \times \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ with each $|\mathcal{S}_i| = 2$, $\mathcal{S}_i = \{a, b\}$. Given T samples from either P^1 or P^2 , we want to tell which distribution the samples are coming from. We will show that, if we can solve the forecast aggregation problem, then we can distinguish the two distributions (with high probability), which requires $T = \Omega(\frac{1}{d_H^2(P^1, P^2)} \log \frac{1}{\delta}) = \Omega(\frac{1}{\varepsilon} \log \frac{1}{\delta})$ samples according to Lemma A.5.

Let $c = 32$. We assume $\varepsilon < 2^{-18}$, so that $c\sqrt{\varepsilon} < \frac{1}{16}$. For P^1 , we let

$$P^1(\omega = 1) = 0.5 - \frac{1}{16n} + \frac{c\sqrt{\varepsilon}}{n} =: p^1.$$

For P^2 , we let

$$P^2(\omega = 1) = 0.5 - \frac{1}{16n} - \frac{c\sqrt{\varepsilon}}{n} =: p^2.$$

We require that, in the forecast aggregation problem under both distributions P^1 and P^2 , whenever expert i sees signal a, b , she reports

$$r_a = 0.5, \quad r_b = 0,$$

respectively. This gives the following conditional probabilities $P^1(\cdot | \omega), P^2(\cdot | \omega)$:

$$\begin{bmatrix} P^1(a | \omega = 0) \\ P^1(a | \omega = 1) \end{bmatrix} = \frac{p^1 - r_b}{r_a - r_b} \begin{bmatrix} \frac{1-r_a}{1-p^1} \\ \frac{r_a}{p^1} \end{bmatrix} = \begin{bmatrix} \frac{1-\frac{1}{8n} + \frac{2c\sqrt{\varepsilon}}{n}}{1+\frac{1}{8n} - \frac{2c\sqrt{\varepsilon}}{n}} \\ 1 \end{bmatrix}, \quad \begin{bmatrix} P^1(b | \omega = 0) \\ P^1(b | \omega = 1) \end{bmatrix} = \begin{bmatrix} \frac{\frac{1}{4n} - \frac{4c\sqrt{\varepsilon}}{n}}{1+\frac{1}{8n} - \frac{2c\sqrt{\varepsilon}}{n}} \\ 0 \end{bmatrix}. \quad (34)$$

$$\begin{bmatrix} P^2(a | \omega = 0) \\ P^2(a | \omega = 1) \end{bmatrix} = \frac{p^2 - r_b}{r_a - r_b} \begin{bmatrix} \frac{1-r_a}{1-p^2} \\ \frac{r_a}{p^2} \end{bmatrix} = \begin{bmatrix} \frac{1-\frac{1}{8n} - \frac{2c\sqrt{\varepsilon}}{n}}{1+\frac{1}{8n} + \frac{2c\sqrt{\varepsilon}}{n}} \\ 1 \end{bmatrix}, \quad \begin{bmatrix} P^2(b | \omega = 0) \\ P^2(b | \omega = 1) \end{bmatrix} = \begin{bmatrix} \frac{\frac{1}{4n} + \frac{4c\sqrt{\varepsilon}}{n}}{1+\frac{1}{8n} + \frac{2c\sqrt{\varepsilon}}{n}} \\ 0 \end{bmatrix}. \quad (35)$$

Given T samples from the unknown distribution $P \in \{P^1, P^2\}$, each of which is a vector of $\omega^{(t)}$ and all experts' signals $s_i^{(t)} \in \{a, b\}$, we feed the corresponding reports $r_i^{(t)} \in \{r_a, r_b\}$ and $\omega^{(t)}$ to the forecast aggregation problem and obtain a solution \hat{f} , which is an ε -optimal aggregator. We want to use \hat{f} to estimate the prior $p = P(\omega = 1) \in \{p^1, p^2\}$ so that we can tell apart P^1 and P^2 . Recall from Lemma 5.1 that $f^*(\mathbf{r}) = \frac{1}{1+\rho^{n-1} \prod_{i=1}^n \frac{1-r_i}{r_i}}$, where $\rho = \frac{p}{1-p}$. Writing ρ in terms of $f^*(\mathbf{r})$, we have

$$\rho = \sqrt[n-1]{\left(\frac{1}{f^*(\mathbf{r})} - 1\right) \prod_{i=1}^n \frac{r_i}{1-r_i}}.$$

In particular, when $r_i = r_a = 0.5$ for all $i \in \{1, \dots, n\}$, we have:

$$\rho = \sqrt[n-1]{\frac{1}{f^*(\mathbf{r}_{0.5})} - 1}, \quad \mathbf{r}_{0.5} = (0.5, \dots, 0.5).$$

So, we estimate ρ by:

$$\hat{\rho} = \sqrt[n-1]{\frac{1}{\hat{f}(\mathbf{r}_{0.5})} - 1}.$$

Now, we want to argue that, if \hat{f} is ε -optimal, then $|\hat{\rho} - \rho|$ is at most $O(\frac{\sqrt{\varepsilon}}{n})$. Consider the function:

$$h(x) = \sqrt[n-1]{\frac{1}{x} - 1},$$

whose derivative is

$$h'(x) = -\frac{1}{n-1} \left(\frac{x}{1-x}\right)^{1-\frac{1}{n-1}} \frac{1}{x^2}.$$

By definition, we have

$$|\hat{\rho} - \rho| = |h(\hat{f}(\mathbf{r}_{0.5})) - h(f^*(\mathbf{r}_{0.5}))|. \quad (36)$$

Claim F.1. $\frac{1}{2} \leq f^*(\mathbf{r}_{0.5}) \leq \frac{2}{3}$.

Proof. For $P \in \{P^1, P^2\}$, its $\rho = \frac{p}{1-p}$ satisfies

$$\begin{aligned} 1 \geq \rho \geq \rho^{n-1} \geq \rho^n &\geq \left(\frac{0.5 - \frac{1}{16n} - \frac{\sqrt{\varepsilon}}{cn}}{0.5 + \frac{1}{16n} + \frac{c\sqrt{\varepsilon}}{n}} \right)^n \\ &= \left(\frac{1 - \frac{1}{8n} - \frac{2c\sqrt{\varepsilon}}{n}}{1 + \frac{1}{8n} + \frac{2c\sqrt{\varepsilon}}{n}} \right)^n > \left(1 - 2\left(\frac{1}{8n} + \frac{2c\sqrt{\varepsilon}}{n}\right) \right)^n \geq 1 - 2\left(\frac{1}{8} + 2c\sqrt{\varepsilon}\right) > \frac{1}{2}, \end{aligned} \quad (37)$$

where in the last three transitions we used the inequalities $\frac{1-x}{1+x} > 1 - 2x$ and $(1 - x/n)^n \geq 1 - x$ for $x \in (0, 1)$ and the fact that $c\sqrt{\varepsilon} < \frac{1}{16}$. So,

$$f^*(\mathbf{r}_{0.5}) = \frac{1}{1 + \rho^{n-1}} \in \left[\frac{1}{1+1}, \frac{1}{1+\frac{1}{2}} \right] = \left[\frac{1}{2}, \frac{2}{3} \right],$$

which proves the claim. \square

With Claim F.1, we can without loss of generality assume $\frac{1}{2} \leq \hat{f}(\mathbf{r}_{0.5}) \leq \frac{2}{3}$ as well (otherwise, we can truncate $\hat{f}(\mathbf{r}_{0.5})$ to this range; this only reduces the approximation error $\mathbb{E}[|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|^2]$).

Claim F.2. For $\frac{1}{2} \leq x \leq \frac{2}{3}$, $|h'(x)| \leq \frac{8}{n-1}$.

Proof.

$$|h'(x)| = \frac{1}{n-1} \left(\frac{x}{1-x} \right)^{1-\frac{1}{n-1}} \frac{1}{x^2} \leq \frac{1}{n-1} \left(\frac{\frac{2}{3}}{1-\frac{2}{3}} \right)^{1-\frac{1}{n-1}} \frac{1}{(\frac{1}{2})^2} = \frac{4}{n-1} \cdot 2^{1-\frac{1}{n-1}} \leq \frac{8}{n-1}. \quad \square$$

Claim F.3. If \hat{f} is ε -optimal, then $|\hat{f}(\mathbf{r}_{0.5}) - f^*(\mathbf{r}_{0.5})| < 2\sqrt{\varepsilon}$.

Proof. If \hat{f} is ε -optimal, i.e., $\mathbb{E}[|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|^2] \leq \varepsilon$, then, by Jensen's inequality $\mathbb{E}[X^2] \geq \mathbb{E}[X]^2$, we have

$$\sqrt{\varepsilon} \geq \mathbb{E}[|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|] = \sum_{\mathbf{r}} P(\mathbf{r}) |\hat{f}(\mathbf{r}) - f^*(\mathbf{r})| \geq P(\mathbf{r}_{0.5}) |\hat{f}(\mathbf{r}_{0.5}) - f^*(\mathbf{r}_{0.5})|. \quad (38)$$

For both $P \in \{P^1, P^2\}$, we have

$$\begin{aligned} P(\mathbf{r}_{0.5}) &= p \cdot P(\mathbf{r}_{0.5} \mid \omega = 1) + (1-p) \cdot P(\mathbf{r}_{0.5} \mid \omega = 0) \\ &= p \cdot P(a \mid \omega = 1)^n + (1-p) \cdot P(a \mid \omega = 0)^n \\ &\geq p \cdot 1 + (1-p) \cdot \left(\frac{1 - \frac{1}{8n} - \frac{2c\sqrt{\varepsilon}}{n}}{1 + \frac{1}{8n} + \frac{2c\sqrt{\varepsilon}}{n}} \right)^n \stackrel{\text{by (37)}}{>} \frac{1}{2}, \end{aligned}$$

Plugging $P(\mathbf{r}_{0.5}) > \frac{1}{2}$ into (38), we get $|\hat{f}(\mathbf{r}_{0.5}) - f^*(\mathbf{r}_{0.5})| < 2\sqrt{\varepsilon}$. \square

From (36), Claim F.1, Claim F.2, and Claim F.3, we get

$$|\hat{\rho} - \rho| = |h(\hat{f}(\mathbf{r}_{0.5})) - h(f^*(\mathbf{r}_{0.5}))| \leq \frac{8}{n-1} |\hat{f}(\mathbf{r}_{0.5}) - f^*(\mathbf{r}_{0.5})| < \frac{8}{n-1} \cdot 2\sqrt{\varepsilon} = \frac{16}{n-1} \sqrt{\varepsilon}. \quad (39)$$

Since $p = \frac{\rho}{1+\rho}$ as a function of ρ has a bounded derivative $\frac{\partial p}{\partial \rho} = \frac{1}{(1+\rho)^2} \leq 1$, Equation (39) implies

$$|\hat{p} - p| < \frac{16}{n-1} \sqrt{\varepsilon}$$

if we use $\hat{p} = \frac{\hat{p}}{1+\hat{p}}$ as an estimate of p . This allows us to tell part P^1 and P^2 because the difference between p^1 and p^2 is greater than twice of our estimation error $|\hat{p} - p|$:

$$|p^1 - p^2| = \frac{2c\sqrt{\varepsilon}}{n} = \frac{64\sqrt{\varepsilon}}{n} \geq \frac{64\sqrt{\varepsilon}}{2(n-1)} = \frac{32\sqrt{\varepsilon}}{n-1} > 2|\hat{p} - p|.$$

Therefore, we can tell part P^1 and P^2 by checking whether p^1 or p^2 is closer to \hat{p} .

Finally, we upper bound the squared Hellinger distance between P^1 and P^2 . This will give the sample complexity lower bound we want.

Claim F.4. $d_{\text{H}}^2(P^1, P^2) \leq O(c^2\varepsilon)$.

Proof. For the marginal distributions of ω , P_{ω}^1 and P_{ω}^2 , according to Lemma A.2 and the fact that $1 \geq \frac{P_{\omega}^2(\omega)}{P_{\omega}^1(\omega)} = \frac{1 - \frac{1}{8n} - \frac{2c\sqrt{\varepsilon}}{n}}{1 - \frac{1}{8n} + \frac{2c\sqrt{\varepsilon}}{n}} = 1 - \frac{\frac{4c\sqrt{\varepsilon}}{n}}{1 - \frac{1}{8n} + \frac{2c\sqrt{\varepsilon}}{n}} = 1 - O\left(\frac{c\sqrt{\varepsilon}}{n}\right)$, we have

$$d_{\text{H}}^2(P_{\omega}^1, P_{\omega}^2) \leq O\left(\left(\frac{c\sqrt{\varepsilon}}{n}\right)^2\right) = O\left(\frac{c^2\varepsilon}{n^2}\right). \quad (40)$$

Given $\omega = 0$ or 1 , we consider the conditional distributions of each s_i , $P_{s_i|\omega}^1$ and $P_{s_i|\omega}^2$. For $s_i = a$, we have

$$\begin{aligned} 1 &\geq \frac{P^2(a|\omega)}{P^1(a|\omega)} \geq \frac{1 - \frac{1}{8n} - \frac{2c\sqrt{\varepsilon}}{n}}{1 + \frac{1}{8n} + \frac{2c\sqrt{\varepsilon}}{n}} \cdot \frac{1 + \frac{1}{8n} - \frac{2c\sqrt{\varepsilon}}{n}}{1 - \frac{1}{8n} + \frac{2c\sqrt{\varepsilon}}{n}} \\ &= \frac{1 + \frac{1}{8n} - \frac{2c\sqrt{\varepsilon}}{n}}{1 + \frac{1}{8n} + \frac{2c\sqrt{\varepsilon}}{n}} \cdot \frac{1 - \frac{1}{8n} - \frac{2c\sqrt{\varepsilon}}{n}}{1 - \frac{1}{8n} + \frac{2c\sqrt{\varepsilon}}{n}} \\ &\left(\frac{a-x}{a+x} \geq 1 - \frac{2x}{a}\right) \geq \left(1 - \frac{\frac{4c\sqrt{\varepsilon}}{n}}{1 + \frac{1}{8n}}\right) \cdot \left(1 - \frac{\frac{4c\sqrt{\varepsilon}}{n}}{1 - \frac{1}{8n}}\right) = 1 - O\left(\frac{c\sqrt{\varepsilon}}{n}\right). \end{aligned}$$

For $s_i = b$, we have

$$1 \geq \frac{P^1(b|\omega)}{P^2(b|\omega)} \geq \frac{\frac{1}{4n} - \frac{4c\sqrt{\varepsilon}}{n}}{1 + \frac{1}{8n} - \frac{2c\sqrt{\varepsilon}}{n}} \cdot \frac{1 + \frac{1}{8n} + \frac{2c\sqrt{\varepsilon}}{n}}{\frac{1}{4n} + \frac{4c\sqrt{\varepsilon}}{n}} \geq \frac{\frac{1}{4n} - \frac{4c\sqrt{\varepsilon}}{n}}{\frac{1}{4n} + \frac{4c\sqrt{\varepsilon}}{n}} = \frac{1 - 16c\sqrt{\varepsilon}}{1 + 16c\sqrt{\varepsilon}} = 1 - O(c\sqrt{\varepsilon}).$$

So, $d_{\text{H}}^2(P_{s_i|\omega}^1, P_{s_i|\omega}^2)$ can be upper bounded as follows:

$$\begin{aligned} d_{\text{H}}^2(P_{s_i|\omega}^1, P_{s_i|\omega}^2) &= \frac{1}{2} \left[\left(\sqrt{P^1(a|\omega)} - \sqrt{P^2(a|\omega)} \right)^2 + \left(\sqrt{P^1(b|\omega)} - \sqrt{P^2(b|\omega)} \right)^2 \right] \\ &= \frac{1}{2} \left[P^1(a|\omega) \left(1 - \sqrt{\frac{P^2(a|\omega)}{P^1(a|\omega)}} \right)^2 + P^2(b|\omega) \left(1 - \sqrt{\frac{P^1(b|\omega)}{P^2(b|\omega)}} \right)^2 \right] \\ &\leq \frac{1}{2} \left[1 \cdot \left(1 - \sqrt{1 - O\left(\frac{c\sqrt{\varepsilon}}{n}\right)} \right)^2 + \left(\frac{1}{4n} + \frac{4c\sqrt{\varepsilon}}{n} \right) \cdot \left(1 - \sqrt{1 - O(c\sqrt{\varepsilon})} \right)^2 \right] \\ \text{(since } 1 - \sqrt{1-x} \leq x) &\leq \frac{1}{2} \left[O\left(\frac{c\sqrt{\varepsilon}}{n}\right)^2 + O\left(\frac{1}{n}\right) \cdot O(c\sqrt{\varepsilon})^2 \right] \\ &= O\left(\frac{c^2\varepsilon}{n}\right). \end{aligned} \quad (41)$$

Since $P^1 = P_{\omega}^1 \cdot \prod_{i=1}^n P_{s_i|\omega}^1$ and $P^2 = P_{\omega}^2 \cdot \prod_{i=1}^n P_{s_i|\omega}^2$, we have, by Lemma A.3 and Lemma A.4,

$$\begin{aligned} d_{\text{H}}^2(P^1, P^2) &\leq d_{\text{H}}^2(P_{\omega}^1, P_{\omega}^2) + \max_{\omega \in \{0,1\}} \left\{ d_{\text{H}}^2\left(\prod_{i=1}^n P_{s_i|\omega}^1, \prod_{i=1}^n P_{s_i|\omega}^2\right) \right\} \\ &\leq d_{\text{H}}^2(P_{\omega}^1, P_{\omega}^2) + \max_{\omega \in \{0,1\}} \left\{ n \cdot d_{\text{H}}^2(P_{s_i|\omega}^1, P_{s_i|\omega}^2) \right\} \\ \text{(40) and (41)} &\leq O\left(\frac{c^2\varepsilon}{n^2}\right) + \max_{\omega \in \{0,1\}} \left\{ n \cdot O\left(\frac{c^2\varepsilon}{n}\right) \right\} \\ &= O(c^2\varepsilon). \quad \square \end{aligned}$$

Therefore, according to Lemma A.5, to tell apart P^1 and P^2 with probability at least $1 - \delta$ we need at least

$$T = \Omega\left(\frac{1}{d_{\text{H}}^2(P^1, P^2)} \log \frac{1}{\delta}\right) = \Omega\left(\frac{1}{c^2 \varepsilon} \log \frac{1}{\delta}\right)$$

samples. This concludes the proof.

G Missing Proofs from Section B

G.1 Proof of Theorem B.2

G.1.1 Additional Notations and Lemmas

We introduce some additional notations and lemmas for the proof. Let μ_0 be the expected average report of all experts conditioning on $\omega = 0$:

$$\mu_0 = \frac{1}{n} \sum_{i=1}^n \mathbb{E}[r_i \mid \omega = 0] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{s_i | \omega=0} [P(\omega = 1 \mid s_i) \mid \omega = 0], \quad (42)$$

which is equal to the expected prediction of ω given expert i 's signal s_i where s_i is distributed conditioning on $\omega = 0$, averaged over all experts. Symmetrically, let

$$\mu_1 = \frac{1}{n} \sum_{i=1}^n \mathbb{E}[1 - r_i \mid \omega = 1] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{s_i | \omega=1} [P(\omega = 0 \mid s_i) \mid \omega = 1]. \quad (43)$$

Recall that $p = P(\omega = 1)$.

Fact G.1. $(1 - p)\mu_0 = p\mu_1$.

Proof. For each expert i , by the law of total expectation and the fact that $r_i = P(\omega = 1 \mid s_i)$, we have the following equations:

$$\begin{aligned} (1 - p) \cdot \mathbb{E}[r_i \mid \omega = 0] + p \cdot \mathbb{E}[r_i \mid \omega = 1] &= P(\omega = 0) \cdot \mathbb{E}[r_i \mid \omega = 0] + P(\omega = 1) \cdot \mathbb{E}[r_i \mid \omega = 1] \\ &= \mathbb{E}[r_i] \\ &= \mathbb{E}_{s_i} [P(\omega = 1 \mid s_i)] = P(\omega = 1) = p. \end{aligned}$$

Subtracting p from both sides, we get

$$(1 - p) \cdot \mathbb{E}[r_i \mid \omega = 0] - p \cdot \mathbb{E}[1 - r_i \mid \omega = 1] = 0.$$

Averaging over all experts $i \in \{1, \dots, n\}$, we conclude that

$$(1 - p) \cdot \frac{1}{n} \sum_{i=1}^n \mathbb{E}[r_i \mid \omega = 0] - p \cdot \frac{1}{n} \sum_{i=1}^n \mathbb{E}[1 - r_i \mid \omega = 1] = 0. \quad \square$$

Lemma G.2. *If $(1 - p)\mu_0 = p\mu_1 < \frac{\varepsilon}{2}$, then the averaging aggregator $f_{\text{avg}}(\mathbf{r}) = \frac{1}{n} \sum_{i=1}^n r_i$ is ε -optimal.*

Proof. If $(1 - p)\mu_0 = p\mu_1 < \frac{\varepsilon}{2}$, then the expected loss of f_{avg} is at most

$$\begin{aligned} L_P(f_{\text{avg}}) &= \mathbb{E}[|f_{\text{avg}} - \omega|^2] \\ &= p \mathbb{E}[(f_{\text{avg}}(\mathbf{r}) - 1)^2 \mid \omega = 1] + (1 - p) \mathbb{E}[(f_{\text{avg}}(\mathbf{r}) - 0)^2 \mid \omega = 0] \\ &\leq p \mathbb{E}[1 - f_{\text{avg}}(\mathbf{r}) \mid \omega = 1] + (1 - p) \mathbb{E}[f_{\text{avg}}(\mathbf{r}) \mid \omega = 0] \\ &= p \mathbb{E}\left[1 - \frac{1}{n} \sum_{i=1}^n r_i \mid \omega = 1\right] + (1 - p) \mathbb{E}\left[\frac{1}{n} \sum_{i=1}^n r_i \mid \omega = 0\right] \\ &= p\mu_1 + (1 - p)\mu_0 \\ &< \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon, \end{aligned}$$

which implies that f_{avg} is ε -optimal. □

The following lemma says that $O(\frac{1}{\varepsilon} \log \frac{1}{\delta})$ samples are sufficient to tell whether the mean of a random variable is below ε or above $\frac{\varepsilon}{2}$:

Lemma G.3. *Given $T = \frac{40}{\varepsilon} \log \frac{2}{\delta}$ i.i.d. samples $X^{(1)}, \dots, X^{(T)}$ of a random variable $X \in [0, 1]$ with unknown mean $\mathbb{E}[X] = \mu$, with probability at least $1 - \delta$ we can tell whether $\mu < \varepsilon$ or $\mu \geq \frac{\varepsilon}{2}$. This can be done by checking whether the empirical mean $\hat{\mu} = \frac{1}{T} \sum_{t=1}^T X^{(t)}$ is $< \frac{3}{4}\varepsilon$ or $\geq \frac{3}{4}\varepsilon$.*

Proof. If $\mu \geq \varepsilon$, using the multiplicative version of Chernoff bound we have

$$\Pr \left[\hat{\mu} < \frac{3}{4}\varepsilon \right] \leq \Pr \left[\hat{\mu} < \frac{3}{4}\mu \right] \leq e^{-\frac{(\frac{1}{4})^2 \mu T}{2}} \leq e^{-\frac{\varepsilon T}{32}} \leq \delta.$$

Namely, with probability at least $1 - \delta$, it holds that

$$\hat{\mu} \geq \frac{3}{4}\varepsilon.$$

If $\mu < \varepsilon$, then using the additive version of Chernoff–Hoeffding theorem, we have

$$\begin{aligned} \Pr \left[\hat{\mu} < \mu - \frac{\varepsilon}{4} \right] &\leq e^{-D(\mu - \frac{\varepsilon}{4} || \mu)T}, \\ \Pr \left[\hat{\mu} > \mu + \frac{\varepsilon}{4} \right] &\leq e^{-D(\mu + \frac{\varepsilon}{4} || \mu)T}, \end{aligned}$$

where $D(x||y) = x \ln \frac{x}{y} + (1-x) \ln \frac{1-x}{1-y}$. Using the inequality $D(x||y) \geq \frac{(x-y)^2}{2y}$ for $x \leq y$ and $D(x||y) \geq \frac{(x-y)^2}{2x}$ for $x \geq y$, we obtain:

$$\begin{aligned} \Pr \left[\hat{\mu} < \mu - \frac{\varepsilon}{4} \right] &\leq e^{-\frac{(\frac{\varepsilon}{4})^2}{2\mu}T} \leq e^{-\frac{(\frac{\varepsilon}{4})^2}{2\varepsilon}T} = e^{-\frac{\varepsilon}{32}T} \leq \frac{\delta}{2}, \\ \Pr \left[\hat{\mu} > \mu + \frac{\varepsilon}{4} \right] &\leq e^{-\frac{(\frac{\varepsilon}{4})^2}{2(\mu + \frac{\varepsilon}{4})}T} \leq e^{-\frac{(\frac{\varepsilon}{4})^2}{2(\frac{5\varepsilon}{4})}T} = e^{-\frac{\varepsilon}{40}T} \leq \frac{\delta}{2}. \end{aligned}$$

By a union bound, with probability at least $1 - \delta$, we have

$$|\hat{\mu} - \mu| \leq \frac{\varepsilon}{4}.$$

Combining the case of $\mu \geq \varepsilon$ and $\mu < \varepsilon$, we conclude that: with probability at least $1 - \delta$,

- If $\hat{\mu} < \frac{3}{4}\varepsilon$, then we must have $\mu < \varepsilon$.
- If $\hat{\mu} \geq \frac{3}{4}\varepsilon$, then we have $\mu \geq \varepsilon$ or $\varepsilon > \mu \geq \hat{\mu} - \frac{\varepsilon}{4} \geq \frac{\varepsilon}{2}$. In either case, we have $\mu \geq \frac{\varepsilon}{2}$.

□

The last lemma we will use shows how to estimate the unknown value of $\rho = \frac{p}{1-p} = \frac{P(\omega=1)}{P(\omega=0)}$ with accuracy Δ using $T = O(\frac{1}{n\Delta^2} \log \frac{1}{\delta})$ samples. Notice that, if one simply uses the empirical value $\hat{\rho} = \frac{\sum_{t=1}^T \mathbb{1}\{\omega^{(t)}=1\}}{\sum_{t=1}^T \mathbb{1}\{\omega^{(t)}=0\}}$ to estimate ρ , then by Chernoff bound this needs $T = O(\frac{1}{\Delta^2} \log \frac{1}{\delta})$ samples, which is larger than what we claim by a factor of n . This sub-optimality is because one only uses the $\omega^{(t)}$'s in the samples to estimate ρ , wasting the reports $r_i^{(t)}$'s. By using $r_i^{(t)}$'s to estimate ρ , we can reduce the number of samples by a factor of n . The basic idea is the following: According to Fact G.1, we have $\rho = \frac{p}{1-p} = \frac{\mu_0}{\mu_1} = \frac{\mathbb{E}[\sum_{i=1}^n r_i | \omega=0]}{\mathbb{E}[\sum_{i=1}^n (1-r_i) | \omega=1]}$. The numerator $\mathbb{E}[\sum_{i=1}^n r_i | \omega=0]$ and the denominator $\mathbb{E}[\sum_{i=1}^n (1-r_i) | \omega=1]$ can be estimated from samples of $r_i^{(t)}$'s where $\omega^{(t)} = 0$ and 1 respectively. The total number of $r_i^{(t)}$'s is Tn , because we have n experts per sample. This reduces the needed number of samples by a factor of n . Formally:

Lemma G.4. *For conditionally independent distribution P , we can estimate $\rho = \frac{P(\omega=1)}{P(\omega=0)}$ with accuracy $\frac{|\hat{\rho}-\rho|}{\rho} \leq \Delta < 1$ (equivalently, $\frac{\hat{\rho}}{\rho} \in 1 \pm \Delta$) with probability at least $1 - \delta$ using*

$$T = O\left(\frac{1}{(1-p)\mu_0 \cdot n \cdot \Delta^2} \log \frac{1}{\delta} + \frac{1}{\min\{p, 1-p\}} \log \frac{1}{\delta} \right)$$

samples of $(\omega^{(t)}, \mathbf{r}^{(t)})$'s, by letting $\hat{\rho} = \frac{\frac{1}{\#_0} \sum_{t:\omega^{(t)}=0} \sum_{i=1}^n r_i^{(t)}}{\frac{1}{\#_1} \sum_{t:\omega^{(t)}=1} \sum_{i=1}^n (1-r_i^{(t)})}$, where $\#_0$ and $\#_1$ are the numbers of samples with $\omega^{(t)} = 0$ and 1 respectively.

Proof. Recall that $p = P(\omega = 1)$, $1 - p = P(\omega = 0)$, and $\rho = \frac{p}{1-p}$. According to Fact G.1 (which says $(1-p)\mu_0 = p\mu_1$), we have

$$\rho = \frac{p}{1-p} = \frac{\mu_0}{\mu_1} = \frac{\sum_{i=1}^n \mathbb{E}[r_i | \omega = 0]}{\sum_{i=1}^n \mathbb{E}[1 - r_i | \omega = 1]}. \quad (44)$$

Consider the following way of estimating ρ from T samples $(\omega^{(t)}, \mathbf{r}^{(t)})_{t=1}^T$: Let $\#_0, \#_1$ be the numbers of samples where $\omega^{(t)} = 0, 1$, respectively:

$$\#_0 = \sum_{t=1}^T \mathbb{1}\{\omega^{(t)} = 0\}, \quad \#_1 = \sum_{t=1}^T \mathbb{1}\{\omega^{(t)} = 1\}.$$

We let

$$\hat{\rho} = \frac{\frac{1}{\#_0} \sum_{t:\omega^{(t)}=0} \sum_{i=1}^n r_i^{(t)}}{\frac{1}{\#_1} \sum_{t:\omega^{(t)}=1} \sum_{i=1}^n (1 - r_i^{(t)})}. \quad (45)$$

Now, we compare the $\hat{\rho}$ in (45) and the ρ in (44): we see that $\frac{1}{\#_0} \sum_{t:\omega^{(t)}=0} \sum_{i=1}^n r_i^{(t)}$ is an (unbiased) estimate of the numerator $\sum_{i=1}^n \mathbb{E}[r_i | \omega = 0] = n\mu_0$ and that $\frac{1}{\#_1} \sum_{t:\omega^{(t)}=1} \sum_{i=1}^n (1 - r_i^{(t)})$ is an (unbiased) estimate of the denominator $\sum_{i=1}^n \mathbb{E}[1 - r_i | \omega = 1] = n\mu_1$. We use Chernoff bounds to argue that the accuracy of the two estimates is within Δ with high probability if $\#_0$ and $\#_1$ are big enough. Suppose that, when drawing the T samples, we draw all the $\omega^{(t)}$'s first (and hence $\#_0, \#_1$ are determined), and then draw all the $r_i^{(t)}$'s. After all the $\omega^{(t)}$'s are drawn, the $r_i^{(t)}$'s become independent, because the signals $s_i^{(t)}$'s are conditionally independent given $\omega^{(t)}$. Therefore, we can use Chernoff bounds:

$$\Pr \left[\left| \frac{1}{\#_0} \sum_{t:\omega^{(t)}=0} \sum_{i=1}^n r_i^{(t)} - n\mu_0 \right| > \Delta n\mu_0 \right] \leq 2e^{-\frac{\#_0 n\mu_0 \Delta^2}{3}},$$

$$\Pr \left[\left| \frac{1}{\#_1} \sum_{t:\omega^{(t)}=1} \sum_{i=1}^n (1 - r_i^{(t)}) - n\mu_1 \right| > \Delta n\mu_1 \right] \leq 2e^{-\frac{\#_1 n\mu_1 \Delta^2}{3}}.$$

Requiring $\delta \geq 2e^{-\frac{\#_0 n\mu_0 \Delta^2}{3}}$ and $\delta \geq 2e^{-\frac{\#_1 n\mu_1 \Delta^2}{3}}$, namely,

$$\#_0 \geq \frac{3}{n\mu_0 \Delta^2} \log \frac{2}{\delta}, \quad \#_1 \geq \frac{3}{n\mu_1 \Delta^2} \log \frac{2}{\delta}, \quad (46)$$

we have, with probability at least $1 - 2\delta$, both of the following hold:

$$\frac{1}{\#_0} \sum_{t:\omega^{(t)}=0} \sum_{i=1}^n r_i^{(t)} \in (1 \pm \Delta)n\mu_0, \quad \frac{1}{\#_1} \sum_{t:\omega^{(t)}=1} \sum_{i=1}^n (1 - r_i^{(t)}) \in (1 \pm \Delta)n\mu_1, \quad (47)$$

Then, we argue that (46) can be satisfied with high probability if T is large enough. This is again done by a Chernoff bound: since $\mathbb{E}[\#_j] = \mathbb{E}[\sum_{t=1}^T \mathbb{1}\{\omega^{(t)} = j\}] = T \cdot P(\omega = j)$, for $j = 0, 1$, we have

$$\Pr \left[\left| \#_0 - T(1-p) \right| \geq \frac{1}{2}T(1-p) \right] \leq 2e^{-\frac{T(1-p)(\frac{1}{2})^2}{3}}, \quad \Pr \left[\left| \#_1 - Tp \right| \geq \frac{1}{2}Tp \right] \leq 2e^{-\frac{Tp(\frac{1}{2})^2}{3}}. \quad (48)$$

So, if we are given

$$T \geq \frac{12}{\min\{p, 1-p\}} \log \frac{2}{\delta} \quad (49)$$

samples, then we can ensure that with probability at least $1 - 2\delta$, it holds $\#_0 \geq \frac{1}{2}T(1-p)$ and $\#_1 \geq \frac{1}{2}Tp$. Then, in order for (46) to be satisfied, we can let

$$\frac{1}{2}T(1-p) \geq \frac{3}{n\mu_0\Delta^2} \log \frac{2}{\delta}, \quad \frac{1}{2}Tp \geq \frac{3}{n\mu_1\Delta^2} \log \frac{2}{\delta}.$$

This gives

$$T \geq \max \left\{ \frac{6}{(1-p)\mu_0 \cdot n\Delta^2} \log \frac{2}{\delta}, \frac{6}{p\mu_1 \cdot n\Delta^2} \log \frac{2}{\delta} \right\} \stackrel{(1-p)\mu_0 = p\mu_1}{=} \frac{6}{(1-p)\mu_0 \cdot n\Delta^2} \log \frac{2}{\delta}. \quad (50)$$

Both (49) and (50) are satisfied when

$$T \geq \frac{6}{(1-p)\mu_0 \cdot n\Delta^2} \log \frac{2}{\delta} + \frac{12}{\min\{p, 1-p\}} \log \frac{2}{\delta}.$$

To conclude, if we are given $T = \frac{6}{(1-p)\mu_0 \cdot n\Delta^2} \log \frac{2}{\delta} + \frac{12}{\min\{p, 1-p\}} \log \frac{2}{\delta}$ samples, then with probability at least $1 - 4\delta$, (47) holds, which implies

$$\hat{\rho} \in \frac{(1 \pm \Delta)\mu_0}{(1 \pm \Delta)\mu_1} = \frac{(1 \pm \Delta)}{(1 \pm \Delta)}\rho \subseteq (1 \pm 4\Delta)\rho \implies \frac{|\hat{\rho} - \rho|}{\rho} \leq 4\Delta,$$

for $\Delta < \frac{1}{4}$. \square

G.1.2 The Proof

We want to show the $O(\frac{1}{\varepsilon n(\frac{1}{1+\gamma})^2} \log \frac{1}{\delta} + \frac{1}{\varepsilon} \log \frac{1}{\delta})$ sample complexity upper bound for the case where experts have γ -strongly informative signals.

We first use $O(\frac{1}{\varepsilon} \log \frac{1}{\delta})$ samples tell whether $(1-p)\mu_0 = p\mu_1 < \frac{\varepsilon}{2}$ or $(1-p)\mu_0 = p\mu_1 \geq \frac{\varepsilon}{4}$. We note that

$$(1-p)\mu_0 = P(\omega = 0) \cdot \mathbb{E}\left[\frac{1}{n} \sum_{i=1}^n r_i \mid \omega = 0\right] = \mathbb{E}\left[\mathbb{1}\{\omega = 0\} \cdot \frac{1}{n} \sum_{i=1}^n r_i\right],$$

which is the expectation of the random variable $X = \mathbb{1}\{\omega = 0\} \cdot \frac{1}{n} \sum_{i=1}^n r_i$. So, according to Lemma G.3, we can tell whether $(1-p)\mu_0 < \frac{\varepsilon}{2}$ or $\geq \frac{\varepsilon}{4}$ with probability at least $1 - \delta$ using $O(\frac{1}{\varepsilon} \log \frac{1}{\delta})$ samples of X . If $(1-p)\mu_0 = p\mu_1 < \frac{\varepsilon}{2}$, then according to Lemma G.2, the averaging aggregator $f_{avg}(\mathbf{r}) = \frac{1}{n} \sum_{i=1}^n r_i$ is ε -optimal. We hence obtained an ε -optimal aggregator in this case. So, in the following proof, we assume $(1-p)\mu_0 = p\mu_1 \geq \frac{\varepsilon}{4}$.

For each expert i , let $\mathcal{S}_i^1 = \{s_i \in \mathcal{S}_i : \frac{P(s_i|\omega=1)}{P(s_i|\omega=0)} \geq 1 + \gamma\}$ be its set of γ -strongly informative signals that are more likely to be realized under $\omega = 1$ than under $\omega = 0$. Let $\mathcal{S}_i^0 = \mathcal{S}_i \setminus \mathcal{S}_i^1 = \{s_i \in \mathcal{S}_i : \frac{P(s_i|\omega=1)}{P(s_i|\omega=0)} \leq \frac{1}{1+\gamma}\}$ be the set of signals that are more likely to be realized under $\omega = 0$. Since $\frac{r_i}{1-r_i} = \frac{P(s_i|\omega=1)}{P(s_i|\omega=0)}\rho$ by Equation (1), whenever an expert receives a signal in \mathcal{S}_i^1 , its report satisfies

$$\frac{r_i}{1-r_i} \geq (1+\gamma)\rho, \quad \forall s_i \in \mathcal{S}_i^1; \quad (51)$$

and whenever it receives a signal in \mathcal{S}_i^0 , its report satisfies

$$\frac{r_i}{1-r_i} \leq \frac{1}{1+\gamma}\rho, \quad \forall s_i \in \mathcal{S}_i^0. \quad (52)$$

We will use the notation $P(\mathcal{S}_i^u \mid \omega) = P(s_i \in \mathcal{S}_i^u \mid \omega) = \sum_{s_i \in \mathcal{S}_i^u} P(s_i \mid \omega)$, for $u \in \{0, 1\}$. Given a set of n signals s_1, \dots, s_n , one per expert, we let $X^1 = \sum_{i=1}^n \mathbb{1}\{s_i \in \mathcal{S}_i^1\}$ be the total number of signals that belong to the \mathcal{S}_i^1 sets, and similarly let $X^0 = \sum_{i=1}^n \mathbb{1}\{s_i \in \mathcal{S}_i^0\}$. We have $X^0 + X^1 = n$, and by definition,

$$\mathbb{E}[X^1 \mid \omega = 1] = \sum_{i=1}^n P(\mathcal{S}_i^1 \mid \omega = 1) \geq (1+\gamma)P(\mathcal{S}_i^1 \mid \omega = 0) = (1+\gamma)\mathbb{E}[X^1 \mid \omega = 0]. \quad (53)$$

$$\mathbb{E}[X^0 \mid \omega = 0] = \sum_{i=1}^n P(\mathcal{S}_i^0 \mid \omega = 0) \geq (1+\gamma)P(\mathcal{S}_i^0 \mid \omega = 1) = (1+\gamma)\mathbb{E}[X^0 \mid \omega = 1], \quad (54)$$

Claim G.5. At least one of $\mathbb{E}[X^1 | \omega = 1]$ and $\mathbb{E}[X^0 | \omega = 0]$ is $\geq \frac{n}{2}$.

Proof. Suppose on the contrary both $\mathbb{E}[X^1 | \omega = 1]$ and $\mathbb{E}[X^0 | \omega = 0]$ are $< \frac{n}{2}$. Then, from (54) we have

$$\mathbb{E}[X^0 | \omega = 1] \leq \frac{1}{1+\gamma} \mathbb{E}[X^0 | \omega = 0] < \frac{n}{2}.$$

This implies $n = \mathbb{E}[X^0 + X^1 | \omega = 1] < \frac{n}{2} + \frac{n}{2} = n$, a contradiction. \square

Let $u \in \{0, 1\}$ be an index such that

$$\mathbb{E}[X^u | \omega = u] \geq \frac{n}{4}. \quad (55)$$

Claim G.5 guarantees that such a u exists. We construct a ‘‘hypothetical’’ aggregator f_{hypo} that, having access to ρ and $\mathbb{E}[X^u | \omega]$, predicts whether $\omega = 0$ or 1 by counting the number X^u of signals that belong to the \mathcal{S}_i^u sets and comparing it with its expectations under $\omega = 0$ and 1, $\mathbb{E}[X^u | \omega = 0]$ and $\mathbb{E}[X^u | \omega = 1]$, respectively. Specifically, given reports $\mathbf{r} = (r_1, \dots, r_n)$ as input, with corresponding unobserved signals $\mathbf{s} = (s_1, \dots, s_n)$, f_{hypo} does the following:

- (1) If $u = 1$, count how many reports r_i 's satisfy $\frac{r_i}{1-r_i} \geq (1+\gamma)\rho$; If $u = 0$, count how many reports r_i 's satisfy $\frac{r_i}{1-r_i} \leq \frac{1}{1+\gamma}\rho$. According to (51) and (52), this number is exactly equal to the number of signals that belong to the \mathcal{S}_i^u sets, X^u .
- (2) Then, check whether X^u is closer (in terms of absolute difference) to $\mathbb{E}[X^u | \omega = u]$ or $\mathbb{E}[X^u | \omega = 1-u]$. If X^u is closer to $\mathbb{E}[X^u | \omega = u]$, output $f_{\text{hypo}}(\mathbf{r}) = u$; otherwise, output $f_{\text{hypo}}(\mathbf{r}) = 1-u$.

We claim that f_{hypo} is ε -optimal.

Claim G.6. Given $\frac{\gamma}{1+\gamma} \geq 8\sqrt{\frac{2}{n} \log \frac{2}{\varepsilon}}$ and $\mathbb{E}[X^u | \omega = u] \geq \frac{n}{4}$, f_{hypo} is ε -optimal.

Proof. Given either $\omega = 0$ or 1, consider the conditional random draw of signals s_1, \dots, s_n . Because $X^u = \sum_{i=1}^n \mathbb{1}\{s_i \in \mathcal{S}_i^u\}$ and the random variables $\mathbb{1}\{s_i \in \mathcal{S}_i^u\}$, $i = 1, \dots, n$, are $[0, 1]$ -bounded and independent conditioning ω , by Hoeffding's inequality we have

$$\Pr \left[|X^u - \mathbb{E}[X^u | \omega]| \geq a \mid \omega \right] \leq 2e^{-\frac{2a^2}{n}}.$$

Let

$$a = \sqrt{\frac{n}{2} \log \frac{2}{\varepsilon}}. \quad (56)$$

Then with probability at least $1 - 2e^{-\frac{2a^2}{n}} = 1 - \varepsilon$, it holds

$$|X^u - \mathbb{E}[X^u | \omega]| < a. \quad (57)$$

Consider the difference between $\mathbb{E}[X^u | \omega = u]$ and $\mathbb{E}[X^u | \omega = 1-u]$. By (53) and (54), we have

$$\mathbb{E}[X^u | \omega = 1-u] \leq \frac{1}{1+\gamma} \mathbb{E}[X^u | \omega = u] = \left(1 - \frac{\gamma}{1+\gamma}\right) \mathbb{E}[X^u | \omega = u].$$

By the assumption $\mathbb{E}[X^u | \omega = u] \geq \frac{n}{4}$,

$$\mathbb{E}[X^u | \omega = 1-u] \leq \mathbb{E}[X^u | \omega = u] - \frac{\gamma}{1+\gamma} \cdot \frac{n}{4}.$$

By the assumption $\frac{\gamma}{1+\gamma} \geq 8\sqrt{\frac{2}{n} \log \frac{2}{\varepsilon}}$, we have $\frac{\gamma}{1+\gamma} \cdot \frac{n}{4} \geq 8\sqrt{\frac{2}{n} \log \frac{2}{\varepsilon}} \cdot \frac{n}{4} = 4\sqrt{\frac{n}{2} \log \frac{2}{\varepsilon}} = 4a$.

Therefore

$$\mathbb{E}[X^u | \omega = u] - \mathbb{E}[X^u | \omega = 1-u] \geq 4a. \quad (58)$$

Because we already had $|X^u - \mathbb{E}[X^u | \omega]| < a$ (which happened with probability at least $1 - \varepsilon$), if X^u turns out to be closer to $\mathbb{E}[X^u | \omega = u]$ it must be that $\omega = u$; if X^u turns out to be closer to $\mathbb{E}[X^u | \omega = 1-u]$ it must be that $\omega = 1-u$. In either case, our output $f_{\text{hypo}}(\mathbf{r})$ is equal to

ω , having a loss 0. If $|X^u - \mathbb{E}[X^u | \omega]| < a$ did not happen, our loss is at most 1. Therefore, the expected loss of our aggregator f_{hyppo} is at most

$$L_P(f_{\text{hyppo}}) = \mathbb{E}_\omega \left[\mathbb{E} [|f_{\text{hyppo}}(\mathbf{r}) - \omega|^2 | \omega] \right] \leq \mathbb{E}_\omega \left[(1 - \varepsilon) \cdot 0 + \varepsilon \cdot 1 \right] = \varepsilon.$$

Since the expected loss of the optimal aggregator f^* is non-negative, f_{hyppo} is ε -optimal. \square

In the remaining proof, we show how to use samples to learn a “real” aggregator \hat{f} that implements the same functionality as the hypothetical aggregator f_{hyppo} and hence is ε -optimal. We have two learning tasks: First, we need to *estimate* ρ , so that we can implement the step (1) of f_{hyppo} which tells apart $\frac{r_i}{1-r_i} \geq (1 + \gamma)\rho$ and $\frac{r_i}{1-r_i} \leq \frac{1}{1+\gamma}\rho$. Second, we need to *find an index* $u \in \{0, 1\}$ such that $\mathbb{E}[X^u | \omega = u] \geq \frac{n}{4}$ and *estimate* $\mathbb{E}[X^u | \omega = u]$, so that we can implement the step (2) of f_{hyppo} which tells whether X^u is closer to $\mathbb{E}[X^u | \omega = u]$ or $\mathbb{E}[X^u | \omega = 1 - u]$. We show that these two tasks can be achieved using $O(\frac{1}{\varepsilon n (\frac{\gamma}{1+\gamma})^2} \log \frac{1}{\delta} + \frac{1}{\varepsilon} \log \frac{1}{\delta})$ samples, with probability at least $1 - O(\delta)$.

Task 1: estimate ρ , using $T_1 = O(\frac{1}{\varepsilon n (\frac{\gamma}{1+\gamma})^2} \log \frac{1}{\delta} + \frac{1}{\varepsilon} \log \frac{1}{\delta})$ samples. We want to use samples to obtain an estimate $\hat{\rho}$ of ρ such that $\frac{1}{1+\gamma}\rho < \hat{\rho} < (1 + \gamma)\rho$. So, by checking whether $\frac{r_i}{1-r_i} > \hat{\rho}$ or $\frac{r_i}{1-r_i} < \hat{\rho}$ we can tell apart $\frac{r_i}{1-r_i} \geq (1 + \gamma)\rho$ and $\frac{r_i}{1-r_i} \leq \frac{1}{1+\gamma}\rho$. Using Lemma G.4 with $\Delta = \frac{\gamma}{1+\gamma}$, we obtain a $\hat{\rho}$ such that

$$\hat{\rho} \in (1 \pm \Delta)\rho,$$

with probability at least $1 - \delta$ using

$$T_1 = O\left(\frac{1}{(1-p)\mu_0 n \Delta^2} \log \frac{1}{\delta} + \frac{1}{\min\{p, 1-p\}} \log \frac{1}{\delta} \right) \leq O\left(\frac{1}{\varepsilon n (\frac{\gamma}{1+\gamma})^2} \log \frac{1}{\delta} + \frac{1}{\varepsilon} \log \frac{1}{\delta} \right)$$

samples (recall that we have $\min\{p, 1-p\} \geq (1-p)\mu_0 = p\mu_1 \geq \frac{\varepsilon}{4}$). The $\hat{\rho}$ then satisfies

$$\hat{\rho} < \left(1 + \frac{\gamma}{1+\gamma}\right)\rho < (1 + \gamma)\rho \quad \text{and} \quad \hat{\rho} > \left(1 - \frac{\gamma}{1+\gamma}\right)\rho = \frac{1}{1+\gamma}\rho,$$

as desired.

Task 2: find u such that $\mathbb{E}[X^u | \omega = u] \geq \frac{n}{4}$ and estimate $\mathbb{E}[X^u | \omega = u]$, using $T_2 = O(\frac{1}{\varepsilon} \log \frac{1}{\delta})$ samples. First, we show how to use $T_2 = O(\frac{1}{\varepsilon} \log \frac{1}{\delta})$ samples to estimate both $\mathbb{E}[X^0 | \omega = 0]$ and $\mathbb{E}[X^1 | \omega = 1]$ with accuracy $a = \sqrt{\frac{n}{2} \log \frac{2}{\varepsilon}}$. By the same argument as in the proof of Lemma G.4 (Equations 48 and 49), we know that with probability at least $1 - 2\delta$ over the random draws of

$$T_2 \geq \frac{12}{\min\{p, 1-p\}} \log \frac{2}{\delta} \tag{59}$$

samples, the numbers of samples $(\omega^{(t)}, r_1^{(t)}, \dots, r_n^{(t)})$'s where $\omega^{(t)} = 0$ and $\omega^{(t)} = 1$, denoted by $\#_0$ and $\#_1$, must satisfy

$$\#_0 \geq \frac{1}{2}(1-p)T_2, \quad \#_1 \geq \frac{1}{2}pT_2.$$

We consider the samples where $\omega^{(t)} = 0$. There are $\#_0 n$ total number of $r_i^{(t)}$'s. Suppose we have accomplished Task 1. Then, for each $r_i^{(t)}$, we can tell whether the corresponding signal $s_i^{(t)}$ belongs to \mathcal{S}_i^0 by checking whether $\frac{r_i^{(t)}}{1-r_i^{(t)}} < \hat{\rho}$. So, we can exactly compute the total number of such signals in the t -th sample, $X^{0(t)} = \sum_{i=1}^n \mathbb{1}\{s_i^{(t)} \in \mathcal{S}_i^0\}$, whose expected value is $\mathbb{E}[X^0 | \omega = 0]$. Because signals are independent given $\omega^{(t)} = 0$, by Hoeffding's inequality we have

$$\Pr \left[\left| \underbrace{\sum_{t: \omega^{(t)}=0} \sum_{i=1}^n \mathbb{1}\{s_i^{(t)} \in \mathcal{S}_i^0\}}_{X^{0(t)}} - \#_0 \mathbb{E}[X^0 | \omega = 0] \right| \geq \#_0 a \right] \leq 2e^{-\frac{2(\#_0 a)^2}{\#_0 n}} = 2e^{-\frac{2\#_0 a^2}{n}}.$$

Plugging in $a = \sqrt{\frac{n}{2} \log \frac{2}{\varepsilon}}$ and $\#_0 \geq \frac{1}{2}(1-p)T_2$, we get

$$\Pr \left[\left| \frac{1}{\#_0} \sum_{t:\omega^{(t)}=0} X^{0(t)} - \mathbb{E}[X^0 | \omega = 0] \right| \geq a \right] \leq 2e^{-\#_0 \log \frac{2}{\varepsilon}} \leq 2e^{-\frac{1}{2}(1-p)T_2 \log \frac{2}{\varepsilon}}.$$

Similarly, considering the samples where $\omega^{(t)} = 1$, we get

$$\Pr \left[\left| \frac{1}{\#_1} \sum_{t:\omega^{(t)}=1} X^{1(t)} - \mathbb{E}[X^1 | \omega = 1] \right| \geq a \right] \leq 2e^{-\#_1 \log \frac{2}{\varepsilon}} \leq 2e^{-\frac{1}{2}pT_2 \log \frac{2}{\varepsilon}}.$$

Therefore, if we require

$$T_2 \geq \frac{2 \log(2/\delta)}{\min\{p, 1-p\} \log(2/\varepsilon)}, \quad (60)$$

then with probability at least $1 - 2\delta$, both

$$\left| \frac{1}{\#_0} \sum_{t:\omega^{(t)}=0} X^{0(t)} - \mathbb{E}[X^0 | \omega = 0] \right| < a, \quad \left| \frac{1}{\#_1} \sum_{t:\omega^{(t)}=1} X^{1(t)} - \mathbb{E}[X^1 | \omega = 1] \right| < a$$

hold. Namely, $\frac{1}{\#_0} \sum_{t:\omega^{(t)}=0} X^{0(t)}$ and $\frac{1}{\#_1} \sum_{t:\omega^{(t)}=1} X^{1(t)}$ are a -accurate estimates of $\mathbb{E}[X^0 | \omega = 0]$ and $\mathbb{E}[X^1 | \omega = 1]$. Equations (59) and (60) together imply that $T_2 = O(\frac{1}{\min\{p, 1-p\}} \log \frac{2}{\delta}) \leq O(\frac{1}{\varepsilon} \log \frac{1}{\delta})$ samples suffice.

Then, we identify an index $u \in \{0, 1\}$ such that $\mathbb{E}[X^u | \omega = u] \geq \frac{n}{4}$. By Claim G.5, there exists a $v \in \{0, 1\}$ with $\mathbb{E}[X^v | \omega = v] \geq \frac{n}{2}$. Since $\frac{1}{\#_v} \sum_{t:\omega^{(t)}=v} X^{v(t)}$ is an a -accurate estimate of $\mathbb{E}[X^v | \omega = v]$, we must have

$$\frac{1}{\#_v} \sum_{t:\omega^{(t)}=v} X^{v(t)} \geq \mathbb{E}[X^v | \omega = v] - a \geq \frac{n}{2} - a.$$

So, at least one of $u \in \{0, 1\}$ must satisfy $\frac{1}{\#_u} \sum_{t:\omega^{(t)}=u} X^{u(t)} \geq \frac{n}{2} - a$. By picking any such a u , we are guaranteed that $\mathbb{E}[X^u | \omega = u] \geq \frac{1}{\#_u} \sum_{t:\omega^{(t)}=u} X^{u(t)} - a \geq \frac{n}{2} - 2a$. Given the assumption $n \geq 32 \log \frac{2}{\varepsilon}$ in the statement of the theorem, we have

$$\frac{a}{n} = \sqrt{\frac{1}{2n} \log \frac{2}{\varepsilon}} \leq \frac{1}{8}.$$

Hence, $\mathbb{E}[X^u | \omega = u] \geq \frac{n}{2} - 2a \geq \frac{n}{2} - 2(\frac{n}{8}) = \frac{n}{4}$.

Finally, as argued above, an a -accurate estimate of $\mathbb{E}[X^u | \omega = u]$ is given by $\frac{1}{\#_u} \sum_{t:\omega^{(t)}=u} X^{u(t)}$.

Constructing \hat{f} . After accomplishing Tasks 1 and 2 using $T_1 + T_2 = O(\frac{1}{\varepsilon n (\frac{1}{1+\gamma})^2} \log \frac{1}{\delta} + \frac{1}{\varepsilon} \log \frac{1}{\delta})$ samples, we construct a \hat{f} that implements the same functionality as f_{hypo} . Let

$$M = \frac{1}{\#_u} \sum_{t:\omega^{(t)}=u} X^{u(t)} - 2a,$$

where $\frac{1}{\#_u} \sum_{t:\omega^{(t)}=u} X^{u(t)}$ is our estimate of $\mathbb{E}[X^u | \omega = u]$ in Task 2 and $a = \sqrt{\frac{n}{2} \log \frac{2}{\varepsilon}}$.

Claim G.7. $\mathbb{E}[X^u | \omega = u] - a > M > \mathbb{E}[X^u | \omega = 1 - u] + a$.

Proof. Because $\frac{1}{\#_u} \sum_{t:\omega^{(t)}=u} X^{u(t)}$ is an a -accurate estimate of $\mathbb{E}[X^u | \omega = u]$, we have

$$\mathbb{E}[X^u | \omega = u] > \frac{1}{\#_u} \sum_{t:\omega^{(t)}=u} X^{u(t)} - a = M + a.$$

Recall from Equation (58) that $\mathbb{E}[X^u | \omega = 1 - u] \leq \mathbb{E}[X^u | \omega = u] - 4a$. So,

$$\mathbb{E}[X^u | \omega = 1 - u] < \left(\frac{1}{\#_u} \sum_{t:\omega^{(t)}=u} X^{u(t)} + a \right) - 4a = M - a.$$

The above two inequalities prove the claim. \square

Given reports $\mathbf{r} = (r_1, \dots, r_n)$ as input, we let \hat{f} do the following:

- (1) If $u = 1$, count how many reports r_i 's satisfy $\frac{r_i}{1-r_i} > \hat{\rho}$; If $u = 0$, count how many reports r_i 's satisfy $\frac{r_i}{1-r_i} < \hat{\rho}$. Let this number be X ;
- (2) Then, check whether $X > M$ or $X \leq M$. If $X > M$, output $\hat{f}(\mathbf{r}) = u$; otherwise, output $\hat{f}(\mathbf{r}) = 1 - u$.

We argue that \hat{f} implements the same functionality as f_{hypo} : (1) In Task 1 we got $\frac{1}{1+\gamma}\rho < \hat{\rho} < (1+\gamma)\rho$. So, by checking whether $\frac{r_i}{1-r_i} > \hat{\rho}$ or $< \hat{\rho}$ we can exactly tell whether $\frac{r_i}{1-r_i} \geq (1+\gamma)\rho$ or $\leq \frac{1}{1+\gamma}\rho$. Hence, we have $X = X^u$, the number of signals that belong to the \mathcal{S}_i^u sets. (2) Recall from Equation (57) that with probability at least $1 - \varepsilon$, X is a -close to its expectation $\mathbb{E}[X^u \mid \omega]$. Then, according to Claim G.7, $X > M$ implies that X is closer to $\mathbb{E}[X^u \mid \omega = u]$; $X < M$ implies that X is closer to $\mathbb{E}[X^u \mid \omega = 1 - u]$. So, \hat{f} implements both of the two steps in f_{hypo} . Hence, according to Claim G.6, \hat{f} is ε -optimal.

G.2 Proof of Theorem B.3

According to Lemma 5.1, the optimal aggregator is

$$f^*(\mathbf{r}) = \frac{1}{1 + \rho^{n-1} \prod_{i=1}^n \frac{1-r_i}{r_i}}, \quad (61)$$

where $\rho = \frac{p}{1-p}$. We claim that an approximately optimal aggregator can be obtained by first estimating ρ from samples and then use the aggregator with the estimate $\hat{\rho}$:

$$\hat{f}(\mathbf{r}) = \frac{1}{1 + \hat{\rho}^{n-1} \prod_{i=1}^n \frac{1-r_i}{r_i}}. \quad (62)$$

Claim G.8. *If $\frac{|\hat{\rho}-\rho|}{\rho} \leq \frac{2\sqrt{\varepsilon}}{n-1} < \frac{1}{2}$, then the aggregator \hat{f} defined above is ε -optimal.*

Proof. Consider the function $g(\hat{\rho}) = \frac{1}{1 + \hat{\rho}^{n-1} \prod_{i=1}^n \frac{1-r_i}{r_i}}$ (where $\hat{\rho}$ is the variable and r_i 's are constants). We claim that

$$|g'(\hat{\rho})| \leq \frac{n-1}{4\hat{\rho}}. \quad (63)$$

To see this, we note that if $r_i = 0$ for some i then $g(\hat{\rho}) = 0$ and $g'(\hat{\rho}) = 0$. Otherwise, we let $y = \prod_{i=1}^n \frac{1-r_i}{r_i} < +\infty$ and take the derivative with respect to $\hat{\rho}$,

$$g'(\hat{\rho}) = -\frac{1}{(1 + \hat{\rho}^{n-1}y)^2} (n-1)\hat{\rho}^{n-2}y = -(n-1) \frac{(\hat{\rho}^{\frac{n}{2}-1}\sqrt{y})^2}{(1 + \hat{\rho}^{n-1}y)^2} = -(n-1) \frac{1}{\left(\frac{1}{\hat{\rho}^{\frac{n}{2}-1}\sqrt{y}} + \hat{\rho}^{\frac{n}{2}}\sqrt{y}\right)^2}.$$

By the AM-GM inequality $a + b \geq 2\sqrt{ab}$, we get

$$|g'(\hat{\rho})| \leq (n-1) \frac{1}{\left(2\sqrt{\frac{1}{\hat{\rho}^{\frac{n}{2}-1}\sqrt{y}} \cdot \hat{\rho}^{\frac{n}{2}}\sqrt{y}}\right)^2} = (n-1) \frac{1}{4\hat{\rho}},$$

as claimed.

Using (63), we have, for $\hat{\rho} \geq \frac{\rho}{2}$,

$$|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})| = |g(\hat{\rho}) - g(\rho)| \leq \frac{n-1}{4 \min\{\hat{\rho}, \rho\}} \cdot |\hat{\rho} - \rho| \leq \frac{n-1}{2} \cdot \frac{|\hat{\rho} - \rho|}{\rho}. \quad (64)$$

So, to obtain ε -approximation $\mathbb{E}[|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|^2] \leq \varepsilon$, we can require $|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})| \leq \frac{\sqrt{\varepsilon}}{2}$. $\frac{|\hat{\rho}-\rho|}{\rho} \leq \sqrt{\varepsilon}$. This can be satisfied if the error in estimating ρ is at most $\frac{|\hat{\rho}-\rho|}{\rho} \leq \frac{2\sqrt{\varepsilon}}{n-1}$. \square

We then show how to use $O(\frac{\gamma n}{\varepsilon} \log \frac{1}{\delta})$ samples to estimate the value of ρ with $O(\frac{\sqrt{\varepsilon}}{n-1})$ accuracy, which will give us an ε -optimal aggregator according to Claim G.8. Recall from (14) that when signals are γ -weakly informative, the reports always satisfy

$$\frac{1}{1+\gamma}\rho \leq \frac{r_i}{1-r_i} \leq (1+\gamma)\rho. \quad (65)$$

The following observation is the key:

Lemma G.9. *For each expert i , we have*

- $\mathbb{E}\left[\frac{r_i}{1-r_i} \mid \omega = 0\right] = \rho;$
- $\mathbb{E}\left[\frac{1-r_i}{r_i} \mid \omega = 1\right] = \frac{1}{\rho}.$

As corollaries, for k conditionally independent reports r_1, \dots, r_k , we have $\mathbb{E}\left[\prod_{i=1}^k \frac{r_i}{1-r_i} \mid \omega = 0\right] = \rho^k$ and $\mathbb{E}\left[\prod_{i=1}^k \frac{1-r_i}{r_i} \mid \omega = 1\right] = \frac{1}{\rho^k}.$

Proof. Because $\frac{r_i}{1-r_i} = \frac{P(s_i|\omega=1)}{P(s_i|\omega=0)}\rho$ (from (13)), we have

$$\mathbb{E}\left[\frac{r_i}{1-r_i} \mid \omega = 0\right] = \sum_{s_i \in \mathcal{S}_i} P(s_i \mid \omega = 0) \frac{P(s_i \mid \omega = 1)}{P(s_i \mid \omega = 0)} \rho = \sum_{s_i \in \mathcal{S}_i} P(s_i \mid \omega = 1) \rho = \rho.$$

For conditionally independent r_1, \dots, r_k , we have

$$\mathbb{E}\left[\prod_{i=1}^k \frac{r_i}{1-r_i} \mid \omega = 0\right] = \prod_{i=1}^k \mathbb{E}\left[\frac{r_i}{1-r_i} \mid \omega = 0\right] = \rho^k.$$

Similarly, we can prove $\mathbb{E}\left[\frac{1-r_i}{r_i} \mid \omega = 1\right] = \frac{1}{\rho}$ and $\mathbb{E}\left[\prod_{i=1}^k \frac{1-r_i}{r_i} \mid \omega = 1\right] = \frac{1}{\rho^k}.$ \square

Let $\Delta = \frac{\sqrt{\varepsilon}}{3\gamma n}$ and suppose we are given $T = \frac{6e}{\gamma n \Delta^2} \log \frac{2}{\delta} = \frac{54e\gamma n}{\varepsilon} \log \frac{2}{\delta} = O(\frac{\gamma n}{\varepsilon} \log \frac{1}{\delta})$ samples. Suppose when drawing the samples we first draw the events $\omega^{(t)}$'s, and then draw the reports $r_i^{(t)}$'s conditioning on $\omega^{(t)}$ being 0 or 1. After the first step, the numbers of samples with $\omega^{(t)} = 0$ and $\omega^{(t)} = 1$ are determined, which we denote by $\#_0$ and $\#_1$. Since $\#_0 + \#_1 = T$, one of them must be at least $T/2$. We argue that whether $\#_0 \geq \frac{T}{2}$ or $\#_1 \geq \frac{T}{2}$ we can estimate $\rho^{1/\gamma}$ with accuracy 3Δ . (For simplicity, we assume that $1/\gamma$ is an integer.)

- If $\#_0 \geq T/2$, then we consider the $\#_0 n$ reports $r_i^{(t)}$'s in the samples with $\omega^{(t)} = 0$. We divide these $\#_0 n$ reports evenly into $\#_0 n \gamma$ groups, each of size $1/\gamma$, denoted by $G_1, \dots, G_{\#_0 n \gamma}$. Consider the product of $\frac{r_i^{(t)}}{1-r_i^{(t)}}$'s in a group G_j : because $r_i^{(t)}$'s are independent given $\omega = 0$, by Lemma G.9 we have

$$\mathbb{E}\left[\prod_{r_i^{(t)} \in G_j} \frac{r_i^{(t)}}{1-r_i^{(t)}} \mid \omega = 0\right] = \rho^{1/\gamma}.$$

Using (65) and the inequality $(1+\gamma)^{1/\gamma} \leq e$, we have

$$\frac{1}{e}\rho^{1/\gamma} \leq \frac{1}{(1+\gamma)^{1/\gamma}}\rho^{1/\gamma} \leq \prod_{r_i^{(t)} \in G_j} \frac{r_i^{(t)}}{1-r_i^{(t)}} \leq (1+\gamma)^{1/\gamma}\rho^{1/\gamma} \leq e\rho^{1/\gamma}.$$

Let X_j be the random variable $\frac{1}{e\rho^{1/\gamma}} \prod_{r_i^{(t)} \in G_j} \frac{r_i^{(t)}}{1-r_i^{(t)}}$. From the above equation and inequality we have $\mathbb{E}[X_j] = \frac{1}{e}$ and $X_j \in [\frac{1}{e^2}, 1] \subseteq [0, 1]$. So, by Chernoff bound,

$$\Pr\left[\frac{1}{\#_0 n \gamma} \sum_{j=1}^{\#_0 n \gamma} X_j \in (1 \pm \Delta) \frac{1}{e} \mid \omega = 0\right] \geq 1 - 2e^{-\frac{\#_0 n \gamma \Delta^2}{3e}} \geq 1 - 2e^{-\frac{T n \gamma \Delta^2}{6e}} = 1 - \delta,$$

given our choice of T . Multiplying $\frac{1}{\#_0 n \gamma} \sum_{j=1}^{\#_0 n \gamma} X_j$ by $e \rho^{1/\gamma}$, we obtain the following estimate of $\rho^{1/\gamma}$:

$$\hat{\rho}_0^{1/\gamma} := \frac{1}{\#_0 n \gamma} \sum_{j=1}^{\#_0 n \gamma} \prod_{r_i^{(t)} \in G_j} \frac{r_i^{(t)}}{1 - r_i^{(t)}} \in (1 \pm \Delta) \rho^{1/\gamma}.$$

Dividing by $\rho^{1/\gamma}$, we get $(\frac{\hat{\rho}_0}{\rho})^{1/\gamma} \in 1 \pm \Delta$.

- If $\#_1 \geq T/2$, then by considering the $\#_1 n$ reports in the samples with $\omega^{(t)} = 1$, dividing them into $\#_1 n \gamma$ groups of size $1/\gamma$, $H_1, \dots, H_{\#_1 n \gamma}$, and similarly defining random variable $Y_j = \frac{\rho^{1/\gamma}}{e} \prod_{r_i^{(t)} \in H_j} \frac{1 - r_i^{(t)}}{r_i^{(t)}}$, we obtain the following estimate of $(\frac{1}{\rho})^{1/\gamma}$:

$$\left(\frac{1}{\hat{\rho}_1}\right)^{1/\gamma} := \frac{1}{\#_1 n \gamma} \sum_{j=1}^{\#_1 n \gamma} \prod_{r_i^{(t)} \in H_j} \frac{1 - r_i^{(t)}}{r_i^{(t)}} \in (1 \pm \Delta) \left(\frac{1}{\rho}\right)^{1/\gamma}.$$

Multiplying by $\rho^{1/\gamma}$, we get $(\frac{\rho}{\hat{\rho}_1})^{1/\gamma} \in 1 \pm \Delta$. Taking the reciprocal and noticing that $\frac{1}{1 \pm \Delta} \subseteq 1 \pm 3\Delta$ when $\Delta < \frac{1}{3}$, we obtain $(\frac{\hat{\rho}_1}{\rho})^{1/\gamma} \in 1 \pm 3\Delta$.

From the discussion above we obtained an estimate $\hat{\rho} \in \{\hat{\rho}_0, \hat{\rho}_1\}$ of ρ such that $(\frac{\hat{\rho}}{\rho})^{1/\gamma} \in 1 \pm 3\Delta$. Raising to the power of γ , and using the inequality $(1-x)^\gamma \geq 1-x\gamma$ and $(1+x)^\gamma \leq e^{x\gamma} \leq 1+2x\gamma$ for $x\gamma \leq 1$, we get

$$\frac{\hat{\rho}}{\rho} \in (1 \pm 3\Delta)^\gamma \subseteq [1 - 3\Delta\gamma, e^{3\Delta\gamma}] \subseteq [1 - 3\Delta\gamma, 1 + 6\Delta\gamma].$$

In particular, this implies $\frac{|\hat{\rho} - \rho|}{\rho} \leq 6\Delta\gamma = \frac{2\sqrt{\varepsilon}}{n}$. Then, according to Claim G.8, the aggregator \hat{f} defined by $\hat{f}(\mathbf{r}) = \frac{1}{1 + \hat{\rho}^{n-1} \prod_{i=1}^n \frac{1 - r_i}{r_i}}$ is ε -optimal. We hence obtained an ε -optimal aggregator.

H Missing Proofs from Section 7

H.1 Proof of Theorem 7.1

Regard P as a joint distribution over reports $\mathbf{r} = (r_1, \dots, r_n)$ and the state ω , where r_i is sampled by first sampling $s_i \in \mathcal{S}_i$ and then letting $r_i = P(\omega = 1 | s_i)$. Since $|\mathcal{S}_i| = m$, there are at most m different values of r_i that can be sampled, so there are at most $2m^n$ different tuples of (\mathbf{r}, ω) in the support of P . For each such tuple (\mathbf{r}, ω) , consider the empirical probability of this tuple:

$$\hat{P}(\mathbf{r}, \omega) = \frac{1}{T} \sum_{t=1}^T \mathbb{1}[(\mathbf{r}^{(t)}, \omega^{(t)}) = (\mathbf{r}, \omega)].$$

By the Chernoff bound, we have

$$\Pr[|\hat{P}(\mathbf{r}, \omega) - P(\mathbf{r}, \omega)| > \Delta P(\mathbf{r}, \omega)] \leq 2e^{-\frac{TP(\mathbf{r}, \omega)\Delta^2}{3}}.$$

Using a union bound for all the $2m^n$ tuples and the fact that $P(\mathbf{r}, \omega) \geq P(\mathbf{s}, \omega) > \frac{c}{m^n}$ (where $\mathbf{s} \in \mathcal{S}$ are some signals that generate \mathbf{r}), we have

$$|\hat{P}(\mathbf{r}, \omega) - P(\mathbf{r}, \omega)| \leq \Delta P(\mathbf{r}, \omega) \tag{66}$$

holds for all tuples (\mathbf{r}, ω) except with probability at most

$$2m^n \cdot 2e^{-\frac{TP(\mathbf{r}, \omega)\Delta^2}{3}} \leq 4m^n \cdot e^{-\frac{cT\Delta^2}{3m^n}} = \delta$$

if

$$T \geq \frac{3m^n}{c\Delta^2} \log \frac{4m^n}{\delta}. \tag{67}$$

Assuming (66) holds, we consider the “empirical” Bayesian aggregator:

$$\hat{f}(\mathbf{r}) = \frac{\hat{P}(\mathbf{r}, \omega)}{\hat{P}(\mathbf{r})}.$$

Since (66) implies $\hat{P}(\mathbf{r}, \omega) \in (1 \pm \Delta)P(\mathbf{r}, \omega)$ and $\hat{P}(\mathbf{r}) \in (1 \pm \Delta)P(\mathbf{r})$, we have

$$\hat{f}(\mathbf{r}) \geq \frac{1 - \Delta}{1 + \Delta} \cdot \frac{P(\mathbf{r}, \omega)}{P(\mathbf{r})} = \left(1 - \frac{2\Delta}{1 + \Delta}\right) f^*(\mathbf{r}) \geq (1 - 2\Delta) f^*(\mathbf{r})$$

and

$$\hat{f}(\mathbf{r}) \leq \frac{1 + \Delta}{1 - \Delta} \cdot \frac{P(\mathbf{r}, \omega)}{P(\mathbf{r})} = \left(1 + \frac{2\Delta}{1 - \Delta}\right) f^*(\mathbf{r}) \leq (1 + 4\Delta) f^*(\mathbf{r}),$$

if $\Delta < \frac{1}{2}$. Putting these two inequalities together:

$$|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})| \leq 4\Delta f^*(\mathbf{r}).$$

We note that this holds for all possible \mathbf{r} in the support of P . So, the expected approximation error of \hat{f} is at most:

$$\begin{aligned} \mathbb{E}[|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|^2] &= \sum_{\mathbf{r}} P(\mathbf{r}) |\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|^2 \\ &\leq \sum_{\mathbf{r}} P(\mathbf{r}) 16\Delta^2 f^*(\mathbf{r})^2 \\ &= 16\Delta^2 \sum_{\mathbf{r}} P(\mathbf{r}) \left(\frac{P(\mathbf{r}, \omega)}{P(\mathbf{r})}\right)^2 \\ &= 16\Delta^2 \sum_{\mathbf{r}} \frac{P(\mathbf{r}, \omega)^2}{P(\mathbf{r})} \leq 16\Delta^2 \sum_{\mathbf{r}} \frac{P(\mathbf{r})^2}{P(\mathbf{r})} = 16\Delta^2. \end{aligned}$$

Letting $16\Delta^2 = \varepsilon$, namely $\Delta^2 = \frac{\varepsilon}{16}$, we have $\mathbb{E}[|\hat{f}(\mathbf{r}) - f^*(\mathbf{r})|^2] \leq \varepsilon$, so \hat{f} is an ε -optimal aggregator. Plugging $\Delta^2 = \frac{\varepsilon}{16}$ into (67) we obtain the sample complexity

$$\frac{3m^n}{c\Delta^2} \log \frac{m^n}{\delta} = \frac{48m^n}{c\varepsilon} \log \frac{m^n}{\delta} = O\left(\frac{m^n}{c\varepsilon} \left(n \log m + \log \frac{1}{\delta}\right)\right).$$