Differentially Private Distributed Inference

Marios Papachristou
Department of Computer Science, Cornell University, papachristoumarios@cs.cornell.edu,

M. Amin Rahimian
Department of Industrial Engineering, University of Pittsburgh, rahimian@pitt.edu,

How can agents exchange information to learn from each other despite their privacy needs and security concerns? Consider healthcare centers that want to collaborate on a multicenter clinical trial, but are concerned about sharing sensitive patient information. Preserving individual privacy and enabling efficient social learning are both important desiderata, but they seem fundamentally at odds. We attempt to reconcile these desiderata by controlling information leakage using statistical disclosure control methods based on differential privacy (DP). Our agents use log-linear rules to update their belief statistics after communicating with their neighbors. DP randomization of beliefs offers communicating agents with plausible deniability with regard to their private information and is amenable to rigorous performance guarantees for the quality of statistical inference. We consider two information environments: one for distributed maximum likelihood estimation (MLE) given a finite number of private signals available at the start of time and another for online learning from an infinite, intermittent stream of private signals that arrive over time. Noisy information aggregation in the finite case leads to interesting trade-offs between rejecting low-quality states and making sure that all high-quality states are admitted in the algorithm output. The MLE setting has natural applications to binary hypothesis testing that we formalize with relevant statistical guarantees. Our results flesh out the nature of the trade-offs between the quality of the inference, learning accuracy, communication cost, and the level of privacy protection that the agents are afforded. In simulation studies, we perform a differentially private, distributed survival analysis on real-world data from an AIDS Clinical Trials Group (ACTG) study to determine whether new treatments improve over standard care. In addition, we used data from clinical trials in advanced cancer patients to determine whether certain biomedical indices affect patient survival. We show that our methods can achieve privacy-preserving inference with significantly more efficient computations than existing privacy-aware methods based on homomorphic encryption, and at lower error rates compared to first-order differentially private distributed optimization methods.

Key words: differential privacy, multicenter clinical trials, consensus, distributed learning, hypothesis testing

History: This version: 3/2025

1. Introduction

In many information environments and information systems, participating agents have privacy and security needs that prevent them from engaging in information exchange at their socially optimal levels (for optimum collective decision making and social learning), e.g., when bound by client confidentiality, obligated by user privacy laws, or when they need to safeguard their information sources (Acquisti et al., 2015; Bélanger and Crossler, 2011; Bélanger and James, 2020). In other social and business settings, participants may be reluctant to express unpopular opinions or openly deliberate on controversial or contentious issues. Even in the absence of such privacy, legal, security, or safety concerns, strategic and competitive considerations can lead to suboptimal information sharing and heterogeneous data integration problems between agents (Zhao and Ram, 2007). It is therefore reasonable to ask the following research question:

How can agents exchange information for collective inference and learn from each other despite their privacy needs and security concerns?

Differential privacy (DP) is a versatile framework for the design of privacy-preserving algorithms and statistical disclosure control. DP has been applied in large-scale and high-profile cases such as the 2020 US decennial census (Garfinkel, 2022) and the tech industry (Erlingsson, 2014; Apple Differential Privacy Team, 2017). Earlier work also points to the potential of DP in alleviating misaligned incentives and reducing competitive inefficiencies in game-theoretic equilibria (Kearns et al., 2014; Blum et al., 2015). We use DP in the design of iterative update rules that allow participating agents to communicate belief statistics with each other without compromising their private information in the belief exchange.

For example, sharing protected health information (PHI) in multicenter clinical trials requires complex data use agreements (DUAs) to protect confidentiality and manage patient authorizations between centers. Clinical trials are the gold standard for establishing medical evidence; however, when testing treatments for emerging diseases, recruiting patients for clinical trials is challenging, especially if the disease is rare and/or severe and there are not many available treatments (Samstein et al., 2019; Froelicher et al., 2021). One solution is to establish a centralized network of clinical trials that collectively recruit patients into a study. For example, in the case of AIDS, this was achieved through the AIDS Clinical Trials Group (ACTG) network, which required substantial prior investment and central coordination (Hammer et al., 1996; Campbell et al., 2012).

However, in the absence of a centrally coordinated network, multicenter clinical trials require the execution of extensive DUAs between healthcare organizations that slow (or sometimes prevent) collaborations. The sharing of clinical trial data remains a significant challenge due to institutional and regulatory barriers, such as the Federal Policy for the Protection of Human Subjects of 1981

(Common Rule), the Health Insurance Portability and Accountability Act of 1996 (HIPAA), the Genetic Information Nondiscrimination Act of 2008 (GINA) in the United States, and the General Data Protection Regulation (GDPR) in the European Union that restrict the sharing of data in healthcare settings and beyond.

The privacy-preserving distributed inference framework in this paper can alleviate the need for complex DUAs and allow centers to communicate privacy-preserving statistics while complying with patient confidentiality and privacy regulations. In 2024, the US Department of Health & Human Services launched the Advancing Clinical Trial Readiness (ACTR) initiative through the ARPA-H Resilient Systems Office (ARPA-H, 2024). ACTR emphasizes the development of a decentralized and on-demand clinical trial infrastructure, aiming to make clinical trials accessible to 90% of eligible participants within 30 minutes of their home. A critical aspect of this transformation is automating data extraction and synchronization between electronic health records (EHRs) and case report forms (CRFs), as well as standardizing data collection across diverse centers. Using data from real-world clinical trials, we demonstrate the utility of our privacy-preserving distributed inference framework for multicenter studies.

This paper provides new methodologies for distributed privacy-preserving inference to improve the quality of information aggregation for collective decision-making in environments where entities harbor privacy concerns that may otherwise deter their efficient participation. These privacy concerns can arise from community retaliation for leaking individual data or the need to comply with regulations protecting individual and patient/client data privacy.

Similarly to the healthcare sector, in finance, education, and development, organizations, while legally obligated to protect client data, can improve service quality and operational efficiency through responsible information sharing (e.g., to improve loan and financial aid screening systems or college admission outcomes) (Adjerid et al., 2018). Our proposed framework can facilitate these goals by allowing multiple entities, such as healthcare care providers, banks, or universities, to perform distributed inference on their private data collectively, without centralizing access to sensitive information, to achieve greater inclusivity, representativeness, and cost efficiency.

1.1. Main Contributions

In this work, we propose a differentially private distributed inference model for information aggregation among a group of agents endowed with private signals and exchanging beliefs to agree on a set of best alternatives, test a hypothesis, or learn a true state of the world. Our agents are wary of revealing their private signals to each other. The private signals are generated according to the true state of the world, which is unknown but common to all agents. We propose non-Bayesian belief exchange models that protect private signals while achieving collective inference and asymptotic learning.

We characterize the quality of learning subject to DP budget constraints and quantify the trade-off between learning rate and privacy: learning in privacy-critical environments can be sustained at a reasonable cost to communication complexity and accuracy. We show that log-linear belief updates, which are used as opinion pools in the decision analysis literature (Gilardoni and Clayton, 1993; Rufo et al., 2012; Abbas, 2009; Rahimian and Jadbabaie, 2015), can be modified to accommodate strict privacy protections for participating agents at a reasonable cost to social learners and their quality of group decision making.

DP constraints in our algorithms are satisfied by adding noise to belief updates at an appropriate level that protects individual signals according to a privacy budget: the less the privacy budget, the more noise is needed to satisfy the DP requirement. In the online learning setting, where agents have access to an infinite, intermittent stream of signals, the effect of noise vanishes asymptotically. Still, DP noising slows down the belief dynamics, requiring more communication rounds than the non-private baseline to achieve the same accuracy level. Our nonasymptotic (finite-time) analysis allows us to delineate the trade-offs between privacy budget, communication complexity, and error probability in the online setting. When agents have access to only finitely many private signals, they aim to agree on a set of alternatives that best explains their collective observations using a distributed, maximum-likelihood estimation (distributed MLE). In this case, the crux of our analysis is that DP noising makes our distributed MLE algorithm non-deterministic. Hence, we need to repeat the distributed MLE algorithm in K rounds to achieve the desired accuracy level. Our approach for DP-distributed MLE opens up a new design dimension to decide how best to aggregate the outcome of different rounds for a reliable MLE output. We propose two methods for aggregating distributed MLE rounds, each of which offers distinct advantages in terms of communication complexity and flexibility to control error probabilities. Our MLE methods have a natural application in hypothesis testing, which we formalize as a differentially private, distributed hypothesis test at a given significance level.

Belief aggregation to detect the best alternatives or test a hypothesis. Our first method is based on *arithmetic* and *geometric* averaging of beliefs across rounds. Specifically, it uses arithmetic averaging as a way to ensure that all good alternatives make their way to the identified set (no Type II errors) and with a control on the false positive rate (Type I errors). Geometric averaging offers a way to control missed detection rate (Type II errors) while ensuring that no bad alternatives are admitted in the output (no Type I errors).

For example, consider evaluating different alternatives to determine the best course of treatment for a critically ill patient. When different alternatives come with severe side effects, it is important to avoid Type I errors (false positives) because a Type I error would admit a less effective treatment while subjecting the patient to unnecessary risks. Geometric averaging is the suitable choice in this case. Similarly, when conducting a clinical trial to test the efficacy of a new drug, controlling the

Type I error rate (false positives) is critical because a Type I error would mean approving ineffective drugs that incur unnecessary costs and harmful side effects. We will show that geometric averaging has a natural extension for distributed hypothesis testing and can be applied in clinical trials to test the efficacy of new drugs. However, when screening patients for early detection of cancer, minimizing the Type II error rate (false negatives) is more important because Type II errors, in this case, reduce the chance of early detection, which delays treatment and leads to worse outcomes for patients. In cases such as cancer screening, we prefer to identify as many true cases as possible, even if it means admitting some false positives (Type I errors), which are resolved at the cost of additional diagnostic tests; therefore, arithmetic aggregation of beliefs is preferred because it precludes Type II errors and allows us to control the Type I error rate.

Our second aggregation method is based on a *thresholding* technique that tracks the frequencies with which beliefs on different states across rounds exceed two thresholds and uses them in two stages to select MLE states. This algorithm works analogously to the averaging algorithms and relies on the concentration of frequencies to control Type I and Type II errors simultaneously using the two thresholds. The advantage of the thresholding method over the geometric and arithmetic averaging algorithms is that it offers more flexibility to control Type I and Type II error rates individually (rather than precluding one and controlling the other); however, this comes at a cost to runtime and communication complexity (an increased number of exchanges of beliefs between the agents) that we characterize in our theoretical performance bounds.

Simulation study of multicenter clinical trials. We test our algorithms on two data sets from clinical trials. Our first experiment considers distributed hypothesis testing for clinical trial data from an AIDS Clinical Trial Group (ACTG) study (Campbell et al., 2012) to determine the effect of different treatments on the survival of AIDS patients using the proportional hazards model (Cox, 1972). In a second experiment, we work with clinical trial data from advanced cancer patients (Samstein et al., 2019) and study the effect of a cancer index – called the Tumor Mutational Burden – on patient survival. Both of our experiments show that our method can efficiently balance distributed learning and privacy while also being easy to implement with significantly faster runtime than existing encryption-based privacy enhancing technologies (PETs) and federated analytics for multicenter studies (Froelicher et al., 2021).

Comparison with state-of-the-art methods. Our algorithmic developments extend the existing work on distributed estimation and learning (Rahimian and Jadbabaie, 2016b) to incorporate privacy guarantees. Although our present work and Rahimian and Jadbabaie (2016b) both consider log-linear learning dynamics in a discrete parameter space, the dynamics in our work are randomized, requiring a substantially different analysis to understand privacy-communication-accuracy tradeoffs in different information environments. Our adaptation of these tradeoffs to distributed hypothesis

testing in healthcare settings is entirely new and not present in the prior literature. Our results are related to the work of Papachristou and Rahimian (2024), who give DP learning algorithms and convergence guarantees in expectation to learn a continuous parameter. Our analysis differs significantly from the one by Papachristou and Rahimian (2024) because our proposed algorithms accommodate a variety of statistical models for inference in discrete spaces. Additionally, we give non-asymptotic guarantees on error probabilities, which allow us to extend our algorithms to admit the type of statical guarantees required for distributed hypothesis testing. This differs significantly from the analysis of Papachristou and Rahimian (2024) that bounds the expected L_2 norm of the error. We also show that our method can achieve significantly lower error rates than first-order distributed optimization methods with similar privacy guarantees (Rizk et al., 2023), and is significantly faster and more efficient than state-of-the-art distributed survival analysis methods based on homomorphic encryption (Froelicher et al., 2021; Geva et al., 2023). A detailed overview of other related work is deferred to Appendix A.

Organization of the paper. In Section 2 we explain a trade-off between error rates when making decisions based on smaller private data sets or larger collective data sets with added noise to protect privacy and then explain the setup for survival analysis in multicenter clinical trials that we use in the simulation study of Section 5. In Section 3 we present our mathematical formulation and performance analysis framework for differentially private distributed inference. In Section 4, we provide our algorithms for differentially private distributed inference along with theoretical guarantees of their performance and privacy in various information environments. Applications of these methods for distributed hypothesis testing, best treatment identification, and online learning on HIV/AIDS and cancer clinical trial data are presented in Section 5, followed by a discussion and concluding remarks in Section 6. We give an extended review of the related work in Appendix A, followed by a description of existing algorithms for distributed MLE and online learning in the non-private regime in Appendix B. The proof details and convergence analysis for distributed private MLE, hypothesis testing, and online learning are provided in Appendices C to E. Additional mathematical details for calculating the sensitivity of the proportional hazards model and lower bounds on the communication complexity of distributed hypothesis testing are given in Appendices F and G, followed by additional simulation experiments in Appendix H. A final notations table is provided in Appendix I.

2. The Dilemma of Privacy-Preserving Data Sharing

We begin with a toy example that helps flesh out a problem at the heart of privacy-preserving data sharing for collective inference: whether to delegate decisions to a central authority with access to shared privatized/noisy data of all agents or to act alone based on clean, individual data? We consider a collection of n agents who want to perform a hypothesis test of a simple null ($\theta = 0$)

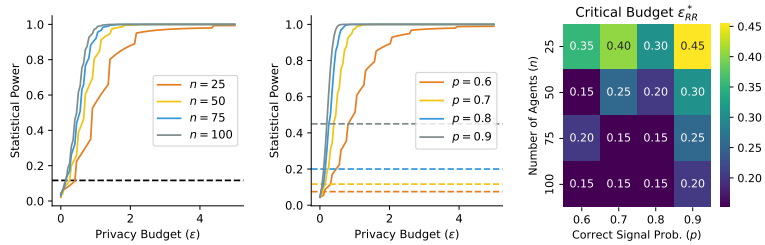

Figure 1 *When inferring a binary state from a binary signal that agrees with the state with probability p, there is a critical privacy budget, $\varepsilon^*_{\mathrm{RR}}$, above which sharing noisy (DP-protected) data becomes beneficial.* **Left:** Statistical power (β_{RR}) as a function of privacy budget ε for different number of agents (n) with $p = 0.7$. The critical budget $\varepsilon^*_{\mathrm{RR}}$ corresponds to the intersection of each curve with the dotted line (β_{IND}) showing the statical power for a single agent. The intersection points show $\varepsilon^*_{\mathrm{RR}}$ decreases with increasing number of agents. **Middle:** Statistical power (β_{RR}) as a function of privacy budget ε for different values of p with $n = 100$ agents. The critical budget ($\varepsilon^*_{\mathrm{RR}}$) corresponds to the intersection points of the power curve for collective hypothesis testing (β_{RR}) with the dotted lines showing statistical power without data sharing for single agents (β_{IND}). Increasing p increases the statical power for both individual agents and collectively when agents share their data; therefore, variation of $\varepsilon^*_{\mathrm{RR}}$ may not be much or monotone. However for very high values pf p where individuals can perform highly accurate tests on their own, the value of $\varepsilon^*_{\mathrm{RR}}$ increases. **Right:** The critical budget ($\varepsilon^*_{\mathrm{RR}}$) values for varying n and p. We use $\alpha = 0.05$ for all tests.

against a simple alternative ($\theta = 1$) at the significance level α; for example, $\alpha = 0.05$. Agents receive binary signals $\boldsymbol{s}_i \in \{0, 1\}$ such that $\mathbb{P}[\boldsymbol{s}_i = 1|\theta = 1] = \mathbb{P}[\boldsymbol{s}_i = 0|\theta = 0] = p$ for some $1/2 < p < 1$; i.e., $\mathbb{P}[\boldsymbol{s}_i = 0|\theta = 1] = \mathbb{P}[\boldsymbol{s}_i = 1|\theta = 0] = 1 - p$. The most powerful individual test at the significance level α varies with p as follows:

1. If $1 - p < \alpha$, then the most powerful test at the significance level α will always reject $\theta = 0$ if $\boldsymbol{s}_i = 1$; and if $\boldsymbol{s}_i = 0$, then it will reject $\theta = 0$ with probability $(\alpha - 1 + p)/p$, giving a total statistical power of $\beta_{\mathrm{IND}} = p + (1 - p)(\alpha - 1 + p)/p$.
2. If $1 - p = \alpha$, then the most powerful test at the significance level α would be to reject $\theta = 0$ if $\boldsymbol{s}_i = 1$ and accept otherwise. The statistical power in this case is given by $\beta_{\mathrm{IND}} = \mathbb{P}[\boldsymbol{s}_i = 1|\theta = 1] = p = 1 - \alpha$.
3. If $1 - p > \alpha$, then the most powerful test at the significance level α will reject $\theta = 0$ with probability $\alpha/(1 - p)$ when $\boldsymbol{s}_i = 1$ and never when $\boldsymbol{s}_i = 0$, giving a statistical power of $\beta_{\mathrm{IND}} = p\alpha/(1 - p)$.

Alternatively, agents can choose to send their private signals to a central authority, which performs the hypothesis test on the collective data on their behalf, in which case they need protection against the privacy ramifications of sharing their signals. Differential privacy with a privacy budget ε limits what can be inferred about private signals from the output of an algorithm $\mathcal{A}$ that is ε-DP because any subset R in the range space of $\mathcal{A}$ must satisfy $\mathbb{P}[\mathcal{A}(\boldsymbol{s}_i = 1) \in R]/\mathbb{P}[\mathcal{A}(\boldsymbol{s}_i = 0) \in R] \leq e^{\varepsilon}$. The randomized response (RR) is a simple mechanism that achieves DP with a privacy budget ε to release

private binary signals (Warner, 1965). According to the randomized response, each agent transmits a signal $\boldsymbol{y}_i$, which is equal to $1-\boldsymbol{s}_i$ with probability $p_\varepsilon = 1/(1+e^\varepsilon)$ where $\varepsilon > 0$ is the privacy budget. By applying the randomized response mechanism before sharing their private signals, agents will have plausible deniability as to what their true private signals are. As $\varepsilon \to 0$, agents randomize their signals with probability 1/2, providing no information to the receiver; on the other hand, as $\varepsilon \to \infty$, agents tend to share their true private signals without randomization noise ($p_\varepsilon \to 0$). To construct the most powerful α level test, agents construct the sum statistic $\sum_{i=1}^n \boldsymbol{y}_i$ and choose a threshold τ_{RR} to reject $\theta = 0$ whenever $\sum_{i=1}^n \boldsymbol{y}_i \geq \tau_{\mathrm{RR}}$. Under the null ($\theta = 0$), the sum statistic is distributed as a binomial random variable with size n and success probability p' where $p' = pp_\varepsilon + (1-p)(1-p_\varepsilon)$. Therefore, agents choose the threshold τ_{RR} to be $\tau_{\mathrm{RR}} = Q_{\mathrm{Bin}(n,p')}(1-\alpha)$, where $Q_{\mathrm{Bin}(n,p')}(1-\alpha)$ is the $(1-\alpha)$-quantile of the $\mathrm{Bin}(n,p')$ distribution. The statistical power of the test, denoted as β_{RR}, is given by (Here we ignore the possibility of randomizing the test outcome if $\mathbb{P}\left[\sum_{i=1}^n \boldsymbol{y}_i \geq \tau_{\mathrm{RR}}+1 \middle| \theta = 0\right] < \alpha$ and $\mathbb{P}\left[\sum_{i=1}^n \boldsymbol{y}_i \geq \tau_{\mathrm{RR}} \middle| \theta = 0\right] > \alpha$, for some positive integer $\tau_{\mathrm{RR}} \leq n$. This case does not arise for the values of $n \geq 25$, p, and $\alpha = 0.05$ that we consider. In general, for $n \geq 25$ the binomial distribution is well approximated by a Gaussian, and such cases, even if they arise, will have minimal effect on the power):

$$\beta_{\mathrm{RR}} = \mathbb{P}\left[\sum_{i=1}^n \boldsymbol{y}_i \geq \tau_{\mathrm{RR}} \middle| \theta = 1\right] = 1 - F_{\mathrm{Bin}(n,1-p')}(\tau_{\mathrm{RR}}) = 1 - F_{\mathrm{Bin}(n,1-p')}\left(Q_{\mathrm{Bin}(n,p')}(1-\alpha)\right). \tag{1}$$

Thus, agents are motivated to share information whenever $\beta_{\mathrm{RR}} \geq \beta_{\mathrm{IND}}$. The *critical privacy budget*, $\varepsilon^*_{\mathrm{RR}}$, above which agents prefer collective hypothesis testing using shared noisy data to individual decision making using their private signals, is given by:

$$\varepsilon^*_{\mathrm{RR}} = \inf_{\varepsilon > 0} \left\{ Q_{\mathrm{Bin}(n,p')}(1-\alpha) \leq Q_{\mathrm{Bin}(n,1-p')}(1-\beta_{\mathrm{IND}}), \text{ where } p' = pp_\varepsilon + (1-p)(1-p_\varepsilon) \right\}. \tag{2}$$

Figure 1 shows the critical budgets $\varepsilon^*_{\mathrm{RR}}$ as a function of n and p: $\varepsilon^*_{\mathrm{RR}}$ decreases with increasing n because the benefit of collective hypothesis testing is greater in larger groups. Our simulations further suggest that $\varepsilon^*_{\mathrm{RR}}$ increases as private signals become more informative with increasing p; with such highly informative private signals, individuals can make accurate inferences without data sharing. Note that in most cases, $\varepsilon^*_{\mathrm{RR}}$ is relatively small ($\varepsilon^*_{\mathrm{RR}} < 1$ in all the cases we have tested), and information sharing is beneficial even for moderate values of n.

Our exploration of the statistical power of the randomized response mechanism gives rise to the following question: Does the randomized response mechanism give a sufficiently powerful statistical test among the class of all DP mechanisms for random binary signals? Awan and Slavković (2018) give a negative answer and provide the mechanism with the maximum statistical power for binary signals. The mechanism of Awan and Slavković (2018) works only in the case of binary signals and cannot be applied to general signal distributions. To address this, Kazan et al. (2023) extend the construction of differentially private hypothesis tests for general distributions by splitting the data

sets and applying the mechanism of Awan and Slavković (2018) on the random variable that counts the rejection of the null hypothesis in subsets of the data.

In general, agents may have different sample sizes, denoted by n_i for each agent i, and can have general distributions. One possibility is for each agent to run the algorithm of Kazan et al. (2023) and share the outcome of their privatized hypothesis tests. The results of the privatized hypothesis tests, weighted as a function of the individual sample sizes and signal likelihoods, can then be combined into an aggregate that is protected by post-processing immunity of DP. However, this prevents consensus on the results of the hypothesis test. In applications such as multicenter clinical trials discussed in Section 2.1, agents need to merge their statistics to test a hypothesis collectively. Our methods allow agents to exchange their statistics and learn from each other towards a consensus on the best alternative, not only for binary hypothesis tests but generally in finite, discrete spaces with more than two alternatives.

In this paper, we devise decentralized algorithms for agents to exchange their log-likelihood ratio statistics and use the Laplace mechanism to privatize the signal log-likelihoods. The Laplace mechanism has several advantages, such as being simple to calculate and implement; it can accommodate DP constraints for general signal distributions and is the mechanism that minimizes the mean squared error bound and convergence time for distributed estimation (Papachristou and Rahimian, 2024; Koufogiannis et al., 2015), and has similar optimality properties for convergence time in our case as well (Theorems 2 and 5 in Section 4). It is also the noise distribution that minimizes the sample complexity of differentially private simple hypothesis tests using noisy clamped log-likelihood ratios (Canonne et al., 2019).

For example, in the case of binary signals, when the agents decide to share data with each other, they calculate their private statistics that correspond to the likelihood ratios and add appropriately chosen Laplace noise, i.e. $\log\left(\frac{\ell(\boldsymbol{s}_i|\theta=1)}{\ell(\boldsymbol{s}_i|\theta=0)}\right)+\boldsymbol{d}_i$, where $\boldsymbol{d}_i \sim \mathrm{Lap}(\Delta/\varepsilon)$ with $\Delta = 2|\log(p)-\log(1-p)|$ being the sensitivity of the log-likelihood ratios to the binary private signals. Then, each agent forms $\sum_{i=1}^{n}\left(\log\left(\frac{\ell(\boldsymbol{s}_i|\theta=1)}{\ell(\boldsymbol{s}_i|\theta=0)}\right)+\boldsymbol{d}_i\right)$ and determines a threshold to reject the null hypothesis at the significance level α. The rejection criteria and the corresponding statistical power β_{Laplace} can be found numerically by simulating the distribution of the test statistics under the null and alternative. Figure 2 shows the statistical power of the test as a function of n and p as well as the critical budgets $\varepsilon^*_{\text{Laplace}}$ above which a collective test based on noisy shared data is more powerful than individual tests based on clean private data. Here we also find $\varepsilon^*_{\text{Laplace}} < 1$ in all our simulations.

2.1. Survival Analysis for Multicenter Clinical Trials

Survival analysis is widely used to examine how patients' survival outcomes (i.e., timing of events such as death, relapse, remission, recovery, etc.) vary in response to specific treatments or other

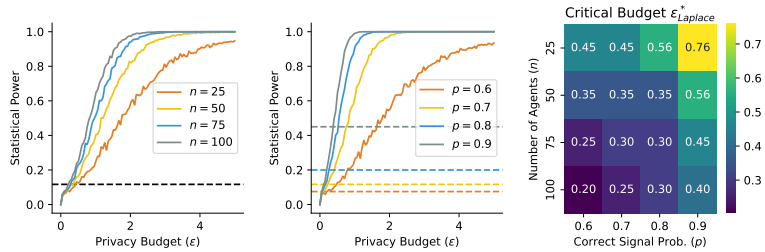

Figure 2 Statistical power versus privacy budget for testing the probability of private Bernoulli random variables. The intersections of the dotted lines (β_{IND}) and the curves (β_{Laplace}) give the critical privacy budget, $\varepsilon^*_{\text{Laplace}}$, above which collective testing using noisy shared data is more powerful than individual tests using private signals. All tests are performed at significance level $\alpha = 0.05$. **Left:** Statistical power with different number of agents (n) and $p = 0.7$. The critical privacy budget $\varepsilon^*_{\text{Laplace}}$ decreases with the increasing number of agents n. **Middle:** Statistical power with different values of p and $n = 100$ agents. $\varepsilon^*_{\text{Laplace}}$ as a function of p. The statistical power for both individual and collective tests increase with p, and therefore the intersection points does not vary monotonically with increasing p; however, $\varepsilon^*_{\text{Laplace}}$ is highest at high values of p where individuals alone can perform highly accurate tests. **Right:** The critical privacy budget $\varepsilon^*_{\text{Laplace}}$ for different values of n and p. The rejection criteria and statistical power are numerically determined by drawing 10,000 samples under the null and alternative ($\theta \in \{0, 1\}$).

health factors. Importantly, survival analysis accounts for censoring, which refers to the existence of data points where events of interest, such as death, have not occurred by the end of the study period, and yet these data points are informative about the effect of treatments. Survival models account for censoring to give more accurate and unbiased estimates. This approach typically involves key metrics such as survival functions, hazard rates, and median survival times, allowing researchers to compare the effectiveness of different treatments or identify prognostic factors. Standard nonparametric and semiparametric models, such as the Kaplan-Meier curve (Kaplan and Meier, 1958) and the Cox proportional hazards model (Cox, 1972), allow us to analyze patient survival data with great flexibility in various clinical and experimental settings. By modeling time-to-event data, survival analysis plays a crucial role in clinical trials, epidemiological studies, and personalized medicine, helping clinicians and researchers make informed decisions about treatment strategies and patient care. Our simulation study in Section 5 consists of survival analysis on a data set from the AIDS Clinical Trials Group (ACTG) network (National Institute of Allergy and Infectious Diseases, 1995), which has conducted several studies of HIV-infected and AIDS patients since the late 1980s (Hammer et al., 1996; Campbell et al., 2012).

Healthcare centers vary in the demographics of the patients that they serve, and they may even have different target groups (e.g., children's or women's hospitals). Although different centers can conduct clinical trials and hypothesis tests on their own, potential data limitations can yield false conclusions, for example, due to small sample sizes or higher prevalence of specific health conditions among some demographics that are under- or overrepresented in their patient population. For this reason, multiple

healthcare centers can join multicenter trials to perform hypothesis tests with more accuracy, using larger and more diverse patient samples. To model multicenter trials, we divide patient data from the ACTG study 175 (National Institute of Allergy and Infectious Diseases, 1995) equally at random between five centers and use the proportional hazards model (Cox, 1972) for survival analysis. For ease of exposition, we consider a simple hypothesis testing scenario to determine the efficacy and safety of a new treatment (ddI) against standard care (ZDV), which are antiretroviral HIV / AIDS medications in the ACTG study 175. For each center $i \in [n]$, we denote their data set of patients by $\boldsymbol{S}_i$ with cardinality n_i. The patient j in the center i has the treatment variable $\boldsymbol{x}_{ij} \in \{0,1\}$ which indicates if they have received the new treatment (ddI), $\boldsymbol{x}_{ij} = 1$, or standard care (ZDV), $\boldsymbol{x}_{ij} = 0$. The effect of the treatment is parametrized by θ. The survival event for each patient is denoted by $\boldsymbol{\delta}_{ij} \in \{0,1\}$ with $\boldsymbol{\delta}_{ij} = 1$ corresponding to death (survival time is observed) and $\boldsymbol{\delta}_{ij} = 0$ corresponding to survival (survival time is censored). The survival or censoring time is measured in days and is indicated by $\boldsymbol{t}_{ij} \in \mathbb{N}$. Following the Cox proportional hazards model, the partial log-likelihood of the patient survival data for center i is given by:

$$\log \ell_i(\boldsymbol{S}_i|\theta) = \sum_{j:\boldsymbol{\delta}_{ij}=1} \left(\boldsymbol{x}_{ij}\theta - \log \sum_{j':\boldsymbol{t}_{ij'} \geq \boldsymbol{t}_{ij}} e^{\boldsymbol{x}_{ij'}\theta} \right). \tag{3}$$

Note that using ℓ_i in (3) is with some abuse of notation because it does not represent the full likelihood of the observed data; however, under Cox's proportional hazards model we can reliably infer the effect of the new treatment and other risk factors that can be included as covariates along with the treatment variable in Equation (3), from the partial likelihoods without specifying the full likelihoods or the functional form of the baseline hazard rate corresponding to $\boldsymbol{x}_{ij} = 0$. We are interested in rejecting the simple null hypothesis that $\theta = 0$, i.e., that the two survival curves are the same, which, in the absence of other covariates, is equivalent to a nonparametric log-rank test of survival times between the treatment and control groups. For concreteness, we use $\theta_1 = -\log(2)$ as the simple alternative, which corresponds to patients being twice as likely to survive under the new treatment than in standard care (e^{θ_1} is the relative risk for the new treatment under the proportional hazards model).

Figure 3(b) shows that the sample size of Kaplan-Meier survival curves calculated by one of the hospitals is not large enough to detect a difference between the two treatments, and, in fact, fitting the proportional hazards model using data from one hospital yields a p-value of 0.9 in Figure 3(c). However, if centers were to pool their data (ignoring privacy concerns) through a centralized authority, then the centralized authority could perform the test at a larger sample size. Figure 3(a) shows the Kaplan-Meier survival curves are easily distinguishable in the centralized case. In Figure 4(c), when all data are pooled together, the difference between log-hazard ratios (log HR) from fitting proportional hazards models for the two treatments is statistically significant at $p < 0.001$.

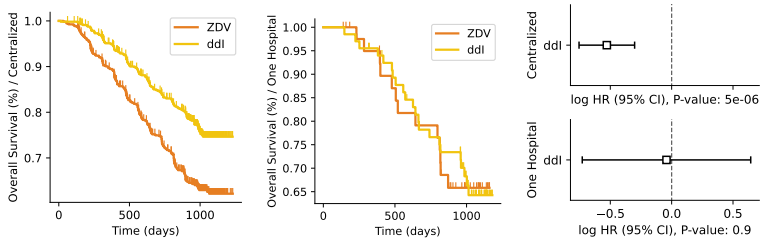

(a) Centralized (b) One hospital (c) Log hazard ratios

Figure 3 **Left:** Kaplan-Meier survival curves for ACTG study 175 (National Institute of Allergy and Infectious Diseases, 1995) for the ZDV and ddI treatments. ZDV stands for Zidovudine, and ddI stands for Didanosine. **Middle:** Survival curves for one hospital (the data is split equally among five hospitals) for the same study. **Right:** Log hazard ratios with 95% confidence intervals from the fitted proportional hazards model using all the data (centralized) and one hospital.

In a distributed setting where privacy concerns limit data sharing and preclude data pooling for centralized access, centers can exchange noisy information locally to achieve distributed hypothesis testing with privacy guarantees. In the private regime, inspired by the likelihood ratio test, we propose that each center calculate its local (partial) log-likelihood ratio statistic $\log\left(\frac{\ell_i(\boldsymbol{S}_i|\theta=\theta_1)}{\ell_i(\boldsymbol{S}_i|\theta=0)}\right)$, add appropriately chosen Laplace noise $\boldsymbol{d}_i$, and then exchange the noisy statistic with its neighbors; see the simulation study in Section 5. The methods we devise in Sections 3 and 4 allow agents to form belief statistics locally, in a privacy-preserving manner, and use them for collective hypothesis testing with guarantees of false positive and false negative rates.

3. Belief Prorogation for Differentially Private Distributed Inference

More generally, each agent can have access to private observations with different likelihoods (for example, representing different populations of patients in different centers), which are parametrized by a common, finite set of alternatives. The goal of the agents is to exchange information to calculate the likelihoods of their collective observations and to choose a common set of maximum likelihood estimates (MLEs). To this end, agents exchange (non-Bayesian) belief statistics and decide on a common set of best alternatives (collective likelihood maximizers) based on their convergent beliefs. This MLE setup has natural extensions to hypothesis testing and online learning that we explore in Section 4. Before focusing on the theoretical performance and privacy guarantees of our distributed belief exchange algorithms in Section 4, we introduce our information environment and the setup of distributed inference in Section 3.1. Our belief exchange rules have a log-linear format, which have normative foundations in decision analysis and can also be justified on the basis of their convergence properties. We discuss these normative foundations and convergence properties in Section 3.2, followed by an explanation of our performance analysis framework and associated metrics in Section 3.3.

3.1. Problem Formulation: Distributed Inference & Learning in Discrete Spaces

A collection of n agents, denoted by the set $[n] = \{1, \ldots, n\}$, wants to select an alternative from a finite set Θ. Each agent i has its own data, which are fixed at the beginning or arrive in an online stream over time. The data follow a parametric model given by its likelihood function $\ell_i(\cdot|\theta), \theta \in \Theta$. The goal of the agents is to exchange information to select a set of alternatives that best describes their data collectively, in a maximum likelihood sense. Information exchange between agents (e.g., collaborating hospitals) is constrained by organizational and legal barriers, as well as *privacy regulations* (e.g., HIPAA or GDPR) that limit who can exchange what information with whom. In our framework, organizational and legal barriers are captured by the structure of the communication network that limits information exchange to permissible local neighborhoods, and protection of private data (e.g., protected health information at each hospital) is achieved through differential privacy.

Agents are connected according to an undirected graph $\mathcal{G}([n], \mathcal{E})$ without self-loops. The neighborhood of agent i in $\mathcal{G}$ is denoted by $\mathcal{N}_i$ and corresponds to all agents with whom agent i can exchange information locally. The graph is associated with an irreducible doubly stochastic adjacency matrix $A = [a_{ij}]_{i,j\in[n]}$ with weights $a_{ij} = 0$ whenever $(i,j) \notin \mathcal{E}$ and weight $a_{ij} > 0$ for all $(i,j) \in \mathcal{E}$, and, moreover, $\sum_{i\in[n]} a_{ij} = \sum_{j\in[n]} a_{ij} = 1$. The n eigenvalues of the adjacency matrix A are ordered by their modulus and denoted by $0 < |\lambda_n(A)| \leq \cdots \leq |\lambda_2(A)| < \lambda_1(A) = 1$, with their associated set of bi-orthonormal eigenvectors, $\{\boldsymbol{l}_i, \boldsymbol{r}_i\}_{i\in[n]}$, satisfying $\|\boldsymbol{l}_i\|_2 = \|\boldsymbol{r}_i\|_2 = \boldsymbol{l}_i^T \boldsymbol{r}_i = 1$ for all $i \in [n]$ and $\boldsymbol{l}_i^T \boldsymbol{r}_j = 0$ for all $i \neq j$. Differential privacy restrictions dictate the protection of individual data, that is, ensuring that information exchange mechanisms between agents satisfy the following criteria for ε-DP.

Definition 1. A mechanism $\mathcal{M}_i : \mathcal{S} \to \mathcal{R}$, acting on private signals, is said to be ε-DP with respect to the signal $\boldsymbol{s} \in \mathcal{S} \subset \mathbb{R}^d$ if and only if for all $X \subseteq \mathcal{R}$ we have that

$$\mathbb{P}[\mathcal{M}_i(\boldsymbol{s}) \in X] \leq e^{\varepsilon} \mathbb{P}[\mathcal{M}_i(\boldsymbol{s}') \in X], \text{ for all } \boldsymbol{s}, \boldsymbol{s}' \in \mathcal{S}, \text{ s.t. } \|\boldsymbol{s} - \boldsymbol{s}'\|_1 \leq 1. \tag{4}$$

The ε-DP constraint on $\mathcal{M}_i$ ensures that as $\varepsilon \to 0$, no information can be inferred about whether $\boldsymbol{s}$ or any of its adjacent points (i.e., $\boldsymbol{s}' \in \mathcal{S}$ such that $\|\boldsymbol{s} - \boldsymbol{s}'\|_1 \leq 1$) are the input that generates the observed output $\mathcal{M}_i(\boldsymbol{s})$. Our primary interest in this work is in belief exchange mechanisms for which the range space $\mathcal{R}$ is the probability simplex of all distributions over the state space Θ, denoted by Simplex(Θ). Defining the appropriate notion of adjacency allows us to control sensitive information leaks. When working with clinical trial data, we declare two observations $\boldsymbol{s}$ and $\boldsymbol{s}'$ adjacent if, and only if, they differ in the data of one patient.

We consider two distributed learning environments: In the first, we estimate the best alternative or perform hypothesis testing in a distributed manner based on a set of initial observations; in the second, we learn the true alternative online by repeatedly observing data streams over time. In the

sequel, we present the two settings with the corresponding nonprivate algorithms in Appendix B, which serve as benchmarks for the DP regime.

Distributed MLE. In the distributed maximum likelihood estimation task (MLE), there is a set of likelihood maximizers $\Theta^\star \subset \Theta$, and agents aim to determine $\Theta^\star$ collectively by combining their private signals while communicating within their local neighborhoods. Specifically, each agent has access to a dataset of n_i i.i.d. observations, $\boldsymbol{S}_i = (\boldsymbol{s}_1, \ldots, \boldsymbol{s}_{n_i})$, where $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_{n_i} \in \mathcal{S}$ obey a statistical model determined by the likelihood function $\ell_i(\boldsymbol{S}_i|\hat{\theta}) = \prod_{j=1}^{n_i} \ell_i(\boldsymbol{s}_j|\hat{\theta}), \quad \hat{\theta} \in \Theta$. The agents' task is, given signals $\boldsymbol{S}_i$ for each agent $i \in [n]$, to find the set of likelihood maximizers, i.e.,

$$\Theta^\star = \arg\max_{\hat{\theta}\in\Theta} \Lambda(\hat{\theta}), \text{ where } \Lambda(\hat{\theta}) = \sum_{i=1}^{n} \log\left(\ell_i(\boldsymbol{S}_i|\hat{\theta})\right). \tag{5}$$

We define $\bar{\Theta} = \Theta \setminus \Theta^\star$ as the set of non-optimal states. We let $f^\star$ represent the proportion of MLE states, i.e., $f^\star = |\Theta^\star|/|\Theta|$. Rahimian and Jadbabaie (2016b) give a non-private algorithm for belief exchange, described in Appendix B.1, that is able to recover $\Theta^\star$ asymptotically.

In the multicenter clinical trial example of Section 2.1, the state space is $\Theta = \{0, \theta_1\}$ where $\theta = 0$ corresponds to the hypothesis that the two treatments are statistically indistinguishable and $\theta = \theta_1 < 0$ corresponds to the hypothesis that ddI improves patient survival compared to ZDV as is the case based on Figure 3; hence, $\Theta^\star = \{\theta_1\}$. The goal of the agents is to collectively infer that $\theta = \theta_1$ is the best state (MLE) given the data of the patients because $\Lambda(\theta_1) > \Lambda(0)$. In Section 4.2, we show how distributed MLE algorithms can be modified for distributed hypothesis testing with statistical significance guarantees.

Online Learning from Intermittent Streams. In the online setting, we consider a network of agents making streams of observations intermittently over time. At each time step t, agent i makes $n_{i,t}$ i.i.d. observations $\boldsymbol{s}_{i,t}^1, \ldots, \boldsymbol{s}_{i,t}^{n_{i,t}}$ that are distributed according to a parametric model $\ell_i(\cdot|\hat{\theta}), \hat{\theta} \in \Theta$; and the number of observations at time t for each agent $i \in [n]$ is independently and identically distributed, $\{n_{i,t}\} \overset{\text{i.i.d.}}{\sim} \mathcal{P}_i$, with mean $\mathbb{E}[n_{i,t}] = \xi_i$ and variance $\mathbb{V}[n_{i,t}] = \chi_i$. In the online setting, we assume that the data are generated according to a true parameter $\theta^\circ \in \Theta$, which is uniquely identifiable from the collective observations of the agents. The collective identifiability condition can be expressed as $\sum_{i=1}^{n} \xi_i D_{KL}\left(\ell_i(\cdot|\bar{\theta})|\ell_i(\cdot|\theta^\circ)\right) > 0, \forall \bar{\theta} \neq \theta^\circ$. Note that, unlike the MLE setting, where agents can increase the power of their inferences by exchanging belief statistics, in the online setting, each agent alone has access to an infinite signal stream.

An informational advantage of collective inference in the online setting is apparent where agents individually face identification problems; for example, for an agent $i \in [n]$, a pair of states ($\hat{\theta}_1$ and $\hat{\theta}_2$) may induce the same distribution of observations, making them undistinguishable from the perspective of agent i. In such cases, by exchanging beliefs, agents can benefit from each other's observational

abilities to collectively identify the truth even though they may face an identification problem individually. Rahimian and Jadbabaie (2016b) give a non-private algorithm for belief exchange that is able to asymptotically recover the true state θ° from a stream of observations (cf. Appendix B.2).

In multicenter clinical trials, the online learning setup corresponds to centers recruiting patients over time. The centers can rely on each other's observational capabilities to recruit a diverse patient demographic, which they may otherwise lack access to and can compromise their ability to detect effects in certain populations.

3.2. Differentially Private Log-Linear Belief Updates and Opinion Pools

The fastest way to (asymptotically) compute the joint log-likelihood would be for the agents to run a linear consensus algorithm (DeGroot, 1974; Olshevsky, 2017) in the (privatized) log-likelihood space for each $\theta \in \Theta$. This rule is equivalent to performing log-linear updates (i.e., geometric means) in a (non-Bayesian) belief space (Rahimian and Jadbabaie, 2016b). On the other hand, simple DeGroot (linear) averaging of beliefs converges to the average of initial likelihoods (rather than log-likelihoods) and can give biased estimates if used for maximum likelihood estimation. Kayaalp et al. (2023) and Lalitha et al. (2018, Corollary 2) point out that log-linear learning converges faster than linear learning in the belief space. Marden and Shamma (2012) provide guarantees on the percentage of time that the action profile will be a potential maximizer in potential games using log-linear best-response dynamics.

The log-linear updates that we use to exchange belief statistics among neighboring agents also have parallels as opinion pools to combine probability distributions from multiple experts in the decision and risk analysis literature (Clemen and Winkler, 1999). Abbas (2009), in particular, proposes scoring rules based on the KL divergence between the expert distributions and aggregate beliefs that yield linear and log-linear aggregates as the solution to the expert aggregation problem: If KL divergence is computed from the aggregate belief to the expert beliefs then we obtain log-linear updates, otherwise (i.e., if the KL divergence is computed from the expert beliefs to the aggregate belief) we obtain linear updates. Denoting expert beliefs by $\boldsymbol{\nu}_j$ with weights $a_{ij}, j \neq i$, the log-linear belief aggregate $\nu^\star$ sets the belief in every state $\boldsymbol{\nu}(\theta)$ proportionally to $\prod_{j\in[n]} \boldsymbol{\nu}_j^{a_{ij}}(\theta)$ and is the solution to $\arg\min \sum_{j\neq i} a_{ij} D_{KL}\left(\boldsymbol{\nu}_i | \boldsymbol{\nu}_j\right)$. Such rules are especially sensitive to experts who put zero or small beliefs on some states and can be useful for rejecting non-MLE states or null hypotheses. In this section, we use the framework of Abbas (2009) with explicitly modeled privacy costs to justify our update rules beyond the algorithmic performance framework of Section 3.3 and our theoretical guarantees in Section 4.

Following Abbas (2009), consider a decision maker (i) who is interested in choosing the best option among a set of alternatives (Θ), given access to their own private information ($\boldsymbol{S}_i$) and other opinions

from n experts. The decision maker forms a time-varying public belief $\boldsymbol{\nu}_{i,t} \in \text{Simplex}(\Theta)$ over Θ for treatments that capture all expert distributions and its own data. Initially (at $t=0$), the agent forms an intrinsic belief $\boldsymbol{\gamma}_i \in \text{Simplex}(\Theta)$ based on its own data, i.e., $\boldsymbol{\gamma}_i(\hat{\theta}) = \frac{\ell_i(\boldsymbol{S}_i|\hat{\theta})}{\sum_{\bar{\theta}\in\Theta} \ell_i(\boldsymbol{S}_i|\bar{\theta})}$ for all $\hat{\theta} \in \Theta$. In subsequent steps ($t \geq 1$), the decision maker chooses an expert at random with probability a_{ij} corresponding to the strength of the connection in the communication matrix A and pays C_i to consult with them and C_i to consult its own data and incurs a revenue R_i regardless of its decision. Since the decision maker consults its own data and communicates its beliefs with others, it must protect its own signals due to privacy risks.

The ε-DP mechanism $\mathcal{M}_i$, satisfying Definition 1, applied to $\boldsymbol{\gamma}_i$ produces $\mathcal{M}_i(\boldsymbol{\gamma}_i)$ which is the privatized belief over the state space $|\Theta|$. Putting everything together, the expected payoff of agent i at time t is given as (see Eq. (9) in Abbas (2009)):

$$U_{i,t}(\boldsymbol{\nu}_{i,t}, \overline{\boldsymbol{\nu}}_{t-1}) = \begin{cases} R_i - C_i \sum_{j=1}^{n} a_{ij} D_{KL}\left(\boldsymbol{\nu}_{i,t} | \boldsymbol{\nu}_{j,t-1}\right) - C_i D_{KL}\left(\boldsymbol{\nu}_{i,t} | \boldsymbol{\nu}_{i,t-1}\right), & t \geq 1 \\ R_i - C_i D_{KL}\left(\boldsymbol{\nu}_{i,0} | \mathcal{M}_i(\boldsymbol{\gamma}_i)\right), & t = 0. \end{cases} \tag{6}$$

To implement the mechanism $\mathcal{M}_i$, the decision maker draws appropriately chosen, without loss of generality, zero-mean noise variables $\boldsymbol{d}_i(\hat{\theta})$ for each state $\hat{\theta} \in \Theta$ and adds them to the corresponding log-likelihoods $\log \boldsymbol{\gamma}_i(\hat{\theta})$. Thus, $\mathcal{M}_i(\boldsymbol{\gamma}_i)$ can be written as the product of the two distributions $\boldsymbol{\gamma}_i$ and $\mathfrak{D}_i$, where $\mathfrak{D}_i(\hat{\theta}) = \frac{e^{\boldsymbol{d}_i(\hat{\theta})}}{\sum_{\bar{\theta}\in\Theta} e^{\boldsymbol{d}_i(\bar{\theta})}}$. If we write $\boldsymbol{\nu}_{i,0}$ as the product of $\boldsymbol{\nu}_{i,0}$ and the uniform distribution $\mathcal{U}_\Theta$ on Θ, then we can use the properties of the KL divergence to write $D_{KL}\left(\boldsymbol{\nu}_{i,0} | \mathcal{M}_i(\boldsymbol{\gamma}_i)\right) = D_{KL}\left(\boldsymbol{\nu}_{i,0} | \boldsymbol{\gamma}_i\right) + D_{KL}\left(\mathcal{U}_\Theta | \mathfrak{D}_i\right)$. The term $D_{KL}\left(\mathcal{U}_\Theta | \mathfrak{D}_i\right)$ corresponds to the privacy loss for the utility in Equation (6) and equals

$$D_{KL}\left(\mathcal{U}_\Theta | \mathfrak{D}_i\right) = \sum_{\hat{\theta}\in\Theta} \frac{1}{|\Theta|} \log\left(\frac{\sum_{\bar{\theta}\in\Theta} e^{\boldsymbol{d}_i(\bar{\theta})}}{|\Theta| e^{\boldsymbol{d}_i(\hat{\theta})}}\right) = \log\left(\sum_{\bar{\theta}\in\Theta} e^{\boldsymbol{d}_i(\bar{\theta})}\right) - \log|\Theta| - \sum_{\hat{\theta}\in\Theta} \boldsymbol{d}_i(\hat{\theta}).$$

Taking the expectation over $\mathcal{M}_i$ we get $\mathbb{E}_{\mathcal{M}_i}\left[D_{KL}\left(\mathcal{U}_\Theta | \mathfrak{D}_i\right)\right] = \mathbb{E}_{\mathcal{M}_i}\left[\log\left(\sum_{\bar{\theta}\in\Theta} e^{\boldsymbol{d}_i(\bar{\theta})}\right)\right] - \log|\Theta|$. Since the KL divergence is nonnegative, the expected privacy loss with respect to $\mathcal{M}_i$ is minimized whenever the noise variables $\boldsymbol{d}_i(\hat{\theta})$ are independent of each other and their distribution does not depend on $\hat{\theta}$, i.e., $\boldsymbol{d}_i(\hat{\theta}) = \boldsymbol{d}_i$.

To satisfy the ε-DP requirement in Definition 1, $\boldsymbol{d}_i$ can be set according to the Laplace mechanism (Dwork and Roth, 2014, Section 3.3), which adds Laplace distribution noise in the log space: $\boldsymbol{d}_i \sim \mathcal{D} = \text{Lap}\left(\frac{|\Theta|\Delta_{n,\Theta}}{\varepsilon}\right)$, where $|\Theta|\Delta_{n,\Theta}/\varepsilon$ is the scale parameter of the Laplace distribution and the global ℓ_1-sensitivity, $\Delta_{n,\Theta}$, is defined as:

$$\Delta_{n,\Theta} = \max_{i\in[n],\hat{\theta}\in\Theta} n_i g_i(\hat{\theta}), \quad \text{where} \quad g_i(\hat{\theta}) = \max_{\boldsymbol{S}_i, \boldsymbol{S}_i' \in \mathcal{S}: \|\boldsymbol{S}_i - \boldsymbol{S}_i'\| \leq 1} \left\| \log \ell_i(\boldsymbol{S}_i|\hat{\theta}) - \log \ell_i(\boldsymbol{S}_i'|\hat{\theta}) \right\|_1. \tag{7}$$

Intuitively, when log-likelihoods are highly sensitive to signal values more noise is required to mask the input signals at the ε-DP level. Here we have used $\varepsilon/|\Theta|$ privacy budget per state to ensure that $\boldsymbol{\nu}_{i,0} \in \text{Simplex}(\Theta)$ is ε-DP overall by the composition property of DP (Dwork and Roth, 2014,

Theorem 3.14), after computing its value at all states $\hat{\theta} \in \Theta$. In Appendix F, we bound (7) for the proportional hazards model. More generally, the signal structure may be such that the sensitivity is unbounded when computed globally over the signal space. In such cases, an additive relaxation of the privacy constraints in Definition 1 can be used with a smoothed notion that looks at sensitivity locally at the realized signal values (Nissim et al., 2007). The Laplace mechanism possesses further benefits since it minimizes the convergence time of the belief exchange algorithms (cf. Section 4 and prior results by Papachristou and Rahimian (2024) and Koufogiannis et al. (2015)). The best update rule that optimizes the time-varying utility of Equation (6) yields the log-linear update rules (Abbas, 2009, Proposition 4). The above can be extended to the online setting where data $\boldsymbol{S}_{i,t}$ with $n_{i,t} \geq 0$ data points for each agent i are intermittently arriving at each iteration t, by introducing a time-varying utility and measuring the likelihood at each iteration t. In that regime, first, the agents form the probability measure $\boldsymbol{\gamma}_{i,t}$ such that $\boldsymbol{\gamma}_{i,t}(\hat{\theta}) = \frac{\ell_i(\boldsymbol{S}_{i,t}|\hat{\theta})}{\sum_{\bar{\theta}\in\Theta} \ell_i(\boldsymbol{S}_{i,t}|\bar{\theta})}$ for all $\hat{\theta} \in \Theta$ (when $n_{i,t} = 0$ we set $\ell_i(\boldsymbol{S}_{i,t}|\hat{\theta}) = 1$ for all $\hat{\theta} \in \Theta$). Then the agents have the time-varying utility:

$$U_{i,t}(\boldsymbol{\nu}_{i,t}, \overline{\boldsymbol{\nu}}_{t-1}) = R_i - C_i \sum_{j=1}^{n} a_{ij} D_{KL}\left(\boldsymbol{\nu}_{i,t} | \boldsymbol{\nu}_{j,t-1}\right) - C_i D_{KL}\left(\boldsymbol{\nu}_{i,t} | \mathcal{M}_i(\boldsymbol{\gamma}_{i,t})\right). \tag{8}$$

Optimizing this utility yields log-linear learning, which we leverage in our OL Algorithm 3.

3.3. Performance Analysis Framework

Our proposed log-linear update rules in Section 3.2 rely on adding zero-mean Laplace noise to the likelihood functions drawn from distributions that are independent of the states. For example, in the MLE problem, in expectation, log-linear updates can identify MLE states $\Theta^\star$. However, the added noise makes the algorithm output random, so the algorithm needs to be repeated to achieve the desired accuracy level. Therefore, we need a systematic way to control the error of the output of any algorithm $\mathcal{A}$. Providing high-probability guarantees for an algorithm $\mathcal{A}$ comes at a cost: The algorithm needs to be repeated in sufficiently many rounds and its outputs combined across the rounds.

Type I and Type II error rates. To control the output uncertainty resulting from privacy-preserving randomization noise, our distributed MLE algorithms, in addition to respecting a set DP budget $\varepsilon > 0$, also admit two types of error guarantees, which we refer to as Type I and Type II error probabilities and denote by α and $1-\beta$, following the hypothesis testing nomenclature. In fact, in Appendix D we show that our distributed MLE algorithms and guarantees can be transformed to conduct a distributed hypothesis test at the significance level α. For a distributed MLE algorithm $\mathcal{A}$ that returns a set of maximum likelihood estimators $\hat{\Theta}^{\mathcal{A}} \subseteq \Theta$, we define the *Type I* error rate α as the probability that $\left\{\Theta^\star \not\supseteq \hat{\Theta}^{\mathcal{A}}\right\}$, which can be controlled by bounding $\mathbb{P}\left[\Theta^\star \supseteq \hat{\Theta}^{\mathcal{A}}\right] \geq 1-\alpha$. Similarly, we define the *Type II* error rate $1-\beta$ as the probability of the event $\left\{\Theta^\star \not\subseteq \hat{\Theta}^{\mathcal{A}}\right\}$, which can be controlled by ensuring that $\mathbb{P}\left[\Theta^\star \subseteq \hat{\Theta}^{\mathcal{A}}\right] \geq \beta$. A Type I error occurs when the algorithm detects

a superset of the MLE set $\Theta^\star$. A Type II error occurs when the algorithm filters too many states, missing some MLE states in its output, thus generating a subset of the MLE set $\Theta^\star$.

Returning to our examples from Section 1.1, when selecting among treatments with severe side effects, we want to limit the possibility of selecting ineffective treatments by ensuring a small Type I error rate α so that $\mathbb{P}\left[\Theta^\star \supseteq \hat{\Theta}^{\mathcal{A}}\right] \geq 1-\alpha$. Similarly, when conducting a clinical trial, we need to ensure efficacy at a desired level of statistical significance (e.g., $\alpha = 0.05$). On the other hand, when developing a new cancer screening tool, we want MLE states to be included in the algorithm output ($\{\Theta^\star \subseteq \hat{\Theta}^{\mathcal{A}}\}$) with high probability, that is, for β to be high (e.g., $\beta = 0.8$). This will ensure that we can detect all patients who are at high risk of cancer ($\Theta^\star$). In this case, it is important to control the Type II error rate so that $\mathbb{P}\left[\Theta^\star \not\subseteq \hat{\Theta}^{\mathcal{A}}\right] \leq 1-\alpha$ for a large enough β, but it may be acceptable to misidentify some patients as high risk (i.e., for $\Theta^{\mathcal{A}}$ to include some non-MLE states). In Section 4.1, we show that arithmetic or geometric averaging of beliefs across the rounds may each be suitable, depending on the type of error that one needs to preclude or control.

Communication complexity. To achieve guarantees on the DP budget and the two error probabilities, we must run our algorithms in $K\left(\varepsilon, \eta, \Theta, \mathcal{G}, \{\ell_i(\cdot|\cdot)\}_{i\in[n]}\right)$ rounds and with $T\left(\varepsilon, \eta, \Theta, \mathcal{G}, \{\ell_i(\cdot|\cdot)\}_{i\in[n]}\right)$ iterations in each round, where η here corresponds to either α or $1-\beta$, or the maximum of the two (depending on the algorithm). We define Communication Complexity $=$ $K\left(\varepsilon, \eta, \Theta, \mathcal{G}, \{\ell_i(\cdot|\cdot)\}_{i\in[n]}\right) \cdot T\left(\varepsilon, \eta, \Theta, \mathcal{G}, \{\ell_i(\cdot|\cdot)\}_{i\in[n]}\right)$ as the total number of belief updates. We summarize our communication complexity bounds for MLE and online learning in Table 1 at the beginning of Section 4 before explaining them in detail. Beyond the privacy budget ε, the error probability η, and the number of agents n, the results of Table 1 also depend on the statistical properties of the information environment and the network structure, as described below:

- Regarding the *private signal structures*, the bounds depend on the largest log-likelihood magnitude in the MLE case:

$$\Gamma_{n,\Theta} = \max_{i\in[n],\hat{\theta}\in\Theta} |\log \boldsymbol{\gamma}_{i,t}(\hat{\theta})|, \tag{9}$$

  the minimum divergence between any non-MLE state and an MLE state (or the true state θ° and any other state $\bar{\theta} \neq \theta^\circ$ in the online learning case):

$$l_{n,\Theta} = \min_{\bar{\theta}\in\bar{\Theta},\theta^\star\in\Theta^\star} \left|\sum_{i=1}^{n} \xi_i D_{KL}\left(\ell_i(\cdot|\bar{\theta})|\ell_i(\cdot|\theta^\star)\right)\right| \text{ for MLE, and } \min_{\bar{\theta}\neq\theta^\circ} \left|\sum_{i=1}^{n} \xi_i D_{KL}\left(\ell_i(\cdot|\bar{\theta})|\ell_i(\cdot|\theta^\circ)\right)\right| \text{ for OL,} \tag{10}$$

  the global sensitivity $\Delta_{n,\Theta} = \max_{i\in[n],t\in\mathbb{N}} \Delta_{i,t}$, the sum of variances of the number of signals $\Xi_n = \sum_{i=1}^{n} \chi_i$, and the maximum deviation of the KL divergence

$$Q_{n,\Theta} = \max_{i\in[n],\bar{\theta}\in\bar{\Theta},\theta^\star\in\Theta^\star} \sqrt{\mathbb{V}\left[\log\left(\frac{\ell_i(\cdot|\bar{\theta})}{\ell_i(\cdot|\theta^\star)}\right)\right]} \text{ for MLE, and } \max_{i\in[n],\bar{\theta}\neq\theta^\circ} \sqrt{\mathbb{V}\left[\log\left(\frac{\ell_i(\cdot|\bar{\theta})}{\ell_i(\cdot|\theta^\circ)}\right)\right]} \text{ for OL.} \tag{11}$$

- Regarding the *graph structure*, our bounds depend on the second-largest eigenvalue modulus (SLEM) of A, $b_n^\star = |\lambda_2(A)|$, and the SLEM of $(A+I)/2$, $a_n^\star = |\lambda_2(A+I)|/2$.

4. Performance and Privacy Guarantees for Distributed Inference

We study distributed belief exchange algorithms to find MLEs with theoretical guarantees of privacy and performance (error probability and communication complexity) in Section 4.1. These non-Bayesian belief statistics have a natural application for distributed hypothesis testing, for which we give statistical significance guarantees in Section 4.2. They can also be used for the online learning of the true alternative based on streaming data, which we analyze in Section 4.3.

Setting	With DP (ours)	Without DP (Rahimian and Jadbabaie, 2016b)
	Communication Complexity with respect to ε, η	
MLE (AM/GM)	$\mathcal{O}_{\lvert\Theta\rvert,n}\left(\log(1/\eta)(\log(1/\eta)+\log(1/\varepsilon)+\log(1/\varrho)\right)$	$\mathcal{O}_{\lvert\Theta\rvert,n}(1)$
MLE (Threshold)	$\mathcal{O}_{\lvert\Theta\rvert,n}\left(\frac{\log(1/\eta)}{\pi^2}(\log(1/\eta)+\log(1/\varepsilon)+\log(1/q))\right)$	$\mathcal{O}_{\lvert\Theta\rvert,n}(1)$
Online Learning	$\mathcal{O}_{\lvert\Theta\rvert,n}\left(\frac{1}{\eta^{3/2}}\frac{1}{\varepsilon}\right)$	$\mathcal{O}_{\lvert\Theta\rvert,n}(1)$
	Communication Complexity with respect to n	
MLE (AM/GM)	$\mathcal{O}_{\lvert\Theta\rvert,\varepsilon,\eta}\left(\max\left\{\log\left(\frac{n}{l_{n,\Theta}}\right), \frac{\log\left(\frac{n(\Gamma_{n,\Theta}+\Delta_{n,\Theta})}{\varrho}\right)}{\log(1/a_n^\star)}\right\}\right)$	$\mathcal{O}_{\lvert\Theta\rvert}\left(\max\left\{\log\left(\frac{n}{l_{n,\Theta}}\right), \frac{\log(n\Gamma_{n,\Theta})}{\log(1/a_n^\star)}\right\}\right)$
MLE (Threshold)	$\mathcal{O}_{\lvert\Theta\rvert,\varepsilon,\eta}\left(\frac{1}{\pi^2}\max\left\{\log\left(\frac{n}{l_{n,\Theta}}\right), \frac{\log\left(\frac{n(\Gamma_{n,\Theta}+\Delta_{n,\Theta})}{q}\right)}{\log(1/a_n^\star)}\right\}\right)$	$\mathcal{O}_{\lvert\Theta\rvert}\left(\max\left\{\log\left(\frac{n}{l_{n,\Theta}}\right), \frac{\log(n\Gamma_{n,\Theta})}{\log(1/a_n^\star)}\right\}\right)$
Online Learning	$\mathcal{O}_{\lvert\Theta\rvert,\varepsilon,\eta}\left(\frac{n(Q_{n,\Theta}+\Delta_{n,\Theta})(\max_{i\in[n]}\xi_i+\sqrt{\Xi_n})}{l_{n,\Theta}(1-b_n^\star)}\right)$	$\mathcal{O}_{\lvert\Theta\rvert}\left(\frac{nQ_{n,\Theta}(\max_{i\in[n]}\xi_i+\sqrt{\Xi_n})}{l_{n,\Theta}(1-b_n^\star)}\right)$
	Communication Complexity with respect to $\lvert\Theta\rvert$	
MLE (AM/GM)	$\mathcal{O}_{n,\varepsilon,\eta}\left(\lvert\Theta\rvert\log\lvert\Theta\rvert\left(\log\lvert\Theta\rvert+\log\left(\frac{\Gamma_{n,\Theta}+\lvert\Theta\rvert\log\lvert\Theta\rvert\Delta_{n,\Theta}}{\varrho}\right)\right)\right)$	$\mathcal{O}_n(\log\Gamma_{n,\Theta})$
MLE (Threshold)	$\mathcal{O}_{n,\varepsilon,\eta}\left(\frac{\log\lvert\Theta\rvert}{\pi^2}\left(\log\lvert\Theta\rvert+\log\left(\frac{\Gamma_{n,\Theta}+\lvert\Theta\rvert\log\lvert\Theta\rvert\Delta_{n,\Theta}}{q}\right)\right)\right)$	$\mathcal{O}_n(\log\Gamma_{n,\Theta})$
Online Learning	$\mathcal{O}_{n,\varepsilon,\eta}\left(\frac{\lvert\Theta\rvert(Q_{n,\Theta}+\lvert\Theta\rvert\Delta_{n,\Theta})}{l_{n,\Theta}}\right)$	$\mathcal{O}_n\left(\frac{\lvert\Theta\rvert Q_{n,\Theta}}{l_{n,\Theta}}\right)$

TABLE 1. Communication complexity of distributed inference algorithms for a fixed privacy budget $\varepsilon > 0$ and target maximum Type I/Type II error $\eta \in (0,1)$. We use $\mathcal{O}_{a_1,\ldots,a_N}(\cdot)$ to denote the big-$\mathcal{O}$ notation where the constants are allowed to depend on $a_1,\ldots,a_N$. The overhead of introducing privacy is outlined in blue. For compactness in notation, η refers to $\max\{\alpha, 1-\beta\}$, ϱ refers to $\min\{\varrho^{\mathrm{AM}}, \varrho^{\mathrm{GM}}\}$, π refers to $\min\{\pi_2, \pi_1\}$, and q refers to $\min\{1-q_2, q_1\}$ (see Theorems 2 and 4). The AM/GM algorithms have matching communication complexities whenever $q = \mathcal{O}(1)$ and $\pi = \mathcal{O}\left(1/\sqrt{\lvert\Theta\rvert}\right)$. The non-DP results are due to Rahimian and Jadbabaie (2016b).

Maximum Likelihood Estimation and Hypothesis Testing. Our first set of results considers the recovery of MLEs $\Theta^\star \subset \Theta$ by using two main types of algorithms: The first set of algorithms averages the beliefs of the K independent rounds resulting from log-linear updates both arithmetically and geometrically. As noted earlier in the paper, arithmetic averaging can recover a superset of $\Theta^\star$ – controlling Type I errors – while geometric averaging can recover a subset of $\Theta^\star$ – controlling Type II errors. The formal description of the algorithm is given in Algorithm 1.

Our second algorithm, instead of averaging beliefs, uses two thresholds to select the states for which the average number of times their belief exceeds a (third) threshold is larger than the two

initial thresholds. The advantage of this algorithm is that it gives more flexibility in controlling Type I and Type II errors, and it can even recover the exact MLE set but at a high cost to runtime. The formal description of the algorithm is given in Algorithm 2. Our main result for the MLE follows: Given Type I and Type II errors α and $1-\beta$, respectively, and $\eta = \max\{\alpha, 1-\beta\}$, the following hold:

- For thresholds $\varrho^{\mathrm{AM}}, \varrho^{\mathrm{GM}} > 0$, $\varrho = \min\{\varrho^{\mathrm{AM}}, \varrho^{\mathrm{GM}}\}$, there is an algorithm that requires
$$K \cdot T = \mathcal{O}\left(|\Theta| \log\left(\frac{|\Theta|}{\eta}\right) \max\left\{\log\left(\frac{n}{l_{n,\Theta}}\right), \frac{\log\left(|\Theta| n(\Gamma_{n,\Theta} + \Delta_{n,\Theta}/\varepsilon)/\varrho\right)}{\log(1/a_n^\star)}\right\}\right)$$
exchanges of the beliefs of an agent and constructs two estimators $\hat{\Theta}_{i,T}^{\mathrm{GM}}$ and $\Theta_{i,T}^{\mathrm{AM}}$ such that $\hat{\Theta}_{i,T}^{\mathrm{GM}} \subseteq \Theta^\star \subseteq \Theta_{i,T}^{\mathrm{AM}}$ with probability at least $1-2\eta$.
- For thresholds $q_1, q_2 \in (0,1)$, accuracy parameters $\pi_1, \pi_2 \in (0,1)$, $q = \min\{1-q_1, q_2\}$, and $\pi = \min\{\pi_1, \pi_2\}$ there is an algorithm that requires
$$K \cdot T = \mathcal{O}\left(\frac{1}{\pi^2} \log\left(\frac{|\Theta|}{\eta}\right) \max\left\{\log\left(\frac{n}{l_{n,\Theta}}\right), \frac{\log\left(|\Theta| n(\Gamma_{n,\Theta} + \Delta_{n,\Theta}/\varepsilon)/q\right)}{\log(1/a_n^\star)}\right\}\right)$$
exchanges of the beliefs of an agent, and constructs two estimators $\hat{\Theta}_{i,T}^{\mathrm{thres},1}$ and $\hat{\Theta}_{i,T}^{\mathrm{thres},2}$ such that $\hat{\Theta}_{i,T}^{\mathrm{thres},1} \subseteq \Theta^\star \subseteq \hat{\Theta}_{i,T}^{\mathrm{thres},2}$ with probability at least $1-2\eta$.
- The noise distributions that optimize the number of belief exchanges of an agent are the Laplace distributions with parameters $\Delta_{n,\Theta}|\Theta|^2 \log(|\Theta|/\eta)/\varepsilon$ (up to constant multiplicative factors).
- The resulting estimates are ε-DP with respect to private signals.
- Theorems 2 and 4 give formal statements of these results.
- The results of the GM algorithm can be extended to design hypothesis tests at significance level α (see Proposition 1).

Online Learning. In the sequel, we provide an algorithm that can learn from intermittent streams of data in an online manner, assuming an identifiable joint model $\Lambda(\cdot)$ and data generated from a distribution with (true) parameter $\theta^\circ \in \Theta$. The algorithm is simple and returns the state with the highest belief value. Our main result for the online learning algorithm follows: Given an accuracy parameter $\eta \in (0,1)$ and the identifiability condition $\sum_{i=1}^n \xi_i D_{KL}\left(\ell_i(\cdot|\bar{\theta}) | \ell_i(\cdot|\theta^\circ)\right) > 0, \forall \bar{\theta} \neq \theta^\circ$:

- There exists an algorithm that requires $K = 1$ round and
$$T = \mathcal{O}\left(\frac{n|\Theta|(Q_{n,\Theta} + \Delta_{n,\Theta}|\Theta|/\varepsilon)\left(\max_{i\in[n]} \xi_i + \sqrt{\frac{\Xi_n}{\eta}}\right)}{l_{n,\Theta}\eta(1-b_n^\star)}\right)$$
exchanges of the beliefs of an agent, and returns an estimator $\hat{\theta}_{i,T}^{\mathrm{OL}}$ such that $\hat{\theta}_{i,T}^{\mathrm{OL}} = \theta^\circ$ with probability at least $1-\eta$. The noise distributions that optimize the number of belief exchanges of an agent are the Laplace distributions with parameters $\Delta_{n,\Theta}|\Theta|/\varepsilon$.
- The resulting estimates are ε-DP with respect to private signals.
- Theorem 5 gives a formal version. Table 1 compares the DP algorithms with their non-DP counterparts and presents the overhead of introducing privacy in each case.

In the sequel, we present our algorithms for differentially private distributed MLE and OL, using log-linear belief updates. Our results are non-asymptotic and determine the upper and lower bounds in the communication complexity $K \cdot T$ for a given privacy budget ε, Type I error rate α, and Type II error rate $1-\beta$.

4.1. Differentially Private Distributed MLE

The first idea behind the MLE algorithm is to introduce multiplicative noise subject to DP to the likelihood functions $\boldsymbol{\gamma}_i(\hat{\theta})$. Adding multiplicative noise – which corresponds to additive noise in the log domain – ensures privacy. The agents then communicate their beliefs about the states with their local neighbors and form estimates of the true MLE over time. However, introducing noise makes the algorithm randomized, and therefore, the agents may misidentify the MLE with a non-trivial probability. To overcome this problem and guarantee Type I (resp. Type II) error at most α (resp. $1-\beta$), we run the algorithm for K independent rounds and combine these K estimates to produce the final beliefs. By setting K accordingly, we can reduce the algorithm's errors, which we show in Theorem 1.

To produce the final estimates, we rely on two ways of aggregating the K rounds of the algorithm: The first approach is to take the arithmetic mean of the K rounds to produce the final belief, which we call the AM estimator. The benefit of the AM estimator is that it can accurately w.h.p. identify all the MLE states, i.e., we can control its Type II error probability: $\mathbb{P}\left[\Theta^\star \not\subseteq \hat{\Theta}^{\mathrm{AM}}_{i,T}\right] \leq 1-\beta$. This estimator has high recall but low precision because it can recover a superset of $\Theta^\star$. The second way to combine the K rounds is to take the geometric mean, which we call the GM estimator. For the GM estimator, we can control the Type I error probability, $\mathbb{P}\left[\hat{\Theta}^{\mathrm{GM}}_{i,T} \not\subseteq \Theta^\star\right] \leq \alpha$, to ensure that it recovers a subset of $\Theta^\star$ with high probability. We first start with a result on the asymptotic behavior of Algorithm 1. Specifically, we show that we can recover the MLE with high probability by repeating the algorithm K times for appropriately chosen K (proved in Appendix C.2):

Theorem 1. *For Algorithm 1, with state-independent noise distributions $\mathcal{D}_i(\hat{\theta};\varepsilon) = \mathcal{D}_i(\varepsilon)$ that satisfy ε-DP and do not depend on the state $\hat{\theta}$; as $\varrho^{\mathrm{AM}} \to \infty$ and $\varrho^{\mathrm{GM}} \to \infty$,*

- *for $K \geq |\bar{\Theta}| \log\left(|\Theta^\star|/(1-\beta)\right)$, we have $\lim_{T\to\infty} \mathbb{P}\left[\Theta^\star \subseteq \hat{\Theta}^{\mathrm{AM}}_{i,T}\right] \geq \beta$ for all $i \in [n]$.*
- *for $K \geq |\Theta^\star| \log\left(|\bar{\Theta}|/\alpha\right)$, we have $\lim_{T\to\infty} \mathbb{P}\left[\hat{\Theta}^{\mathrm{GM}}_{i,T} \subseteq \Theta^\star\right] \geq 1-\alpha$ for all $i \in [n]$.*
- *The beliefs exchanged and the resulting estimates are ε-DP with respect to private signals.*

Repeating Algorithm 1 for $K \geq |\Theta| \log(|\Theta|/\min\{\alpha, 1-\beta\})$ iterations, one can assert $\hat{\Theta}^{\mathrm{GM}}_{i,T} \subseteq \Theta^\star \subseteq \hat{\Theta}^{\mathrm{AM}}_{i,T}$ as $T \to \infty$ with probability $\beta - \alpha$ without knowing $f^\star$. Also, it is interesting to point out that the above theorem holds regardless of the noise, as long as the noise distribution does not depend on the state for a given agent and the noise is ε-DP.

Algorithm 1 Private Distributed MLE (AM/GM)

Inputs: Privacy budget ε, Error probabilities $\alpha, 1-\beta$, Log-belief Thresholds $\varrho^{\mathrm{AM}}, \varrho^{\mathrm{GM}} > 0$.

Initialization: Set the number of iterations T and the number of rounds K as indicated by Theorem 1 (for the asymptotic case) or Theorem 2 (for the non-asymptotic case), and $\tau^{\mathrm{AM}} = 1/(1+e^{\varrho^{\mathrm{AM}}})$ (resp. for τ^{GM}). The DP noise distribution for protecting beliefs is $\mathcal{D}_i(\hat{\theta};\varepsilon)$ which can be set optimally according to Theorem 2.

Procedure: The following is repeated in K rounds indexed by $k \in [K]$. In each round k, agents begin by forming noisy likelihoods $\boldsymbol{\sigma}_{i,k}(\hat{\theta}) = e^{\boldsymbol{d}_{i,k}(\hat{\theta})}\boldsymbol{\gamma}_i(\hat{\theta})$, where $\boldsymbol{\gamma}_i(\hat{\theta}) = \prod_{j=1}^{n_i} \ell_i(\boldsymbol{s}_i^j|\tilde{\theta})$ and $\boldsymbol{d}_{i,k}(\hat{\theta}) \sim \mathcal{D}_i(\hat{\theta};\varepsilon)$ independently across agents $i \in [n]$ and states $\hat{\theta} \in \Theta$. The agents initialize their beliefs to $\boldsymbol{\nu}_{i,k,0}(\hat{\theta}) = \boldsymbol{\sigma}_{i,k}(\hat{\theta})/\sum_{\tilde{\theta}\in\Theta}\boldsymbol{\sigma}_{i,k}(\tilde{\theta})$, and over the next T time steps, they communicate with their neighbors and update their beliefs accordingly:

$$\boldsymbol{\nu}_{i,k,t}(\hat{\theta}) = \frac{\boldsymbol{\nu}_{i,k,t-1}^{1+a_{ii}}(\hat{\theta}) \prod_{j\in\mathcal{N}_i} \boldsymbol{\nu}_{j,k,t-1}^{a_{ij}}(\hat{\theta})}{\sum_{\tilde{\theta}\in\Theta} \boldsymbol{\nu}_{i,k,t-1}^{1+a_{ii}}(\tilde{\theta}) \prod_{j\in\mathcal{N}_i} \boldsymbol{\nu}_{j,k,t-1}^{a_{ij}}(\tilde{\theta})} \quad \text{for every } \hat{\theta}\in\Theta, k\in[K] \text{ and } t\in[T]. \tag{12}$$

After the T iterations, the agents aggregate the outcome of the K rounds:

$$\boldsymbol{\nu}_{i,T}^{\mathrm{AM}}(\hat{\theta}) = \frac{\sum_{k\in[K]}\boldsymbol{\nu}_{i,k,T}(\hat{\theta})}{K}, \quad \boldsymbol{\nu}_{i,T}^{\mathrm{GM}}(\hat{\theta}) = \frac{\prod_{k\in[K]}\boldsymbol{\nu}_{i,k,T}(\hat{\theta})^{1/K}}{\sum_{\tilde{\theta}\in\Theta}\prod_{k\in[K]}\boldsymbol{\nu}_{i,k,T}(\tilde{\theta})^{1/K}} \quad \text{for every } \hat{\theta}\in\Theta. \tag{13}$$

Outputs: Return

$$\hat{\Theta}_{i,T}^{\mathrm{AM}} = \left\{\hat{\theta}\in\Theta : \boldsymbol{\nu}_{i,T}^{\mathrm{AM}}(\hat{\theta}) \geq \tau^{\mathrm{AM}}\right\}, \quad \hat{\Theta}_{i,T}^{\mathrm{GM}} = \left\{\hat{\theta}\in\Theta : \boldsymbol{\nu}_{i,T}^{\mathrm{GM}}(\hat{\theta}) \geq \tau^{\mathrm{GM}}\right\}.$$

The distributed MLE algorithm is repeated K times for each of the $|\Theta|$ states, and an attacker would have more observations of the outputs from the same signals; therefore, we need to rescale the privacy budget by $K|\Theta|$ to account for the number of runs and states. Thus, choosing $\mathcal{D}_i(\varepsilon) = \mathrm{Lap}\left(\frac{K|\Theta|\Delta_{n,\Theta}}{\varepsilon}\right)$ satisfies ε-DP and the assumptions of Theorem 1. Our nonasymptotic analysis (presented next) shows that $\mathcal{D}_i(\varepsilon) = \mathrm{Lap}\left(\frac{K|\Theta|\Delta_{n,\Theta}}{\varepsilon}\right)$ is indeed optimal to minimize the required number of iterations per round.

The previous analysis focused on the regime where $T \to \infty$. The next question that arises is to study the behavior of Algorithm 1 for finite t. As expected, the speed of convergence depends on the sum of standard deviations of the noise, that is, $V_{n,\Theta} = \sum_{i=1}^{n}\sqrt{\mathbb{V}_{\boldsymbol{d}_i\sim\mathcal{D}_i(\varepsilon)}[\boldsymbol{d}_i]}$. Our main result for private distributed MLE using the AM/GM method follows (proved in Appendix C.3):

Theorem 2. *The following holds for Algorithm 1, and any $\varrho^{\mathrm{AM}}, \varrho^{\mathrm{GM}} > 0$:*

- *For $T \geq \max\left\{\frac{\log\left(\frac{2\varrho n}{l_{n,\Theta}}\right)}{\log 2}, \frac{\log\left(\frac{|\Theta|^2(n-1)(n\Gamma_{n,\Theta}+V_{n,\Theta})}{2\alpha\varrho^{\mathrm{GM}}\sqrt{K}}\right)}{\log(1/a_n^\star)}\right\}$ and $K \geq |\Theta^\star|\log(|\bar{\Theta}|/\alpha)$, we have $\mathbb{P}[\hat{\Theta}_{i,T}^{\mathrm{GM}} \subseteq \Theta^\star] \geq 1-2\alpha$.*

- *For* $T \geq \max\left\{\frac{\log\left(\frac{2\varrho n}{l_{n,\Theta}}\right)}{\log 2}, \frac{\log\left(\frac{|\Theta|^2(n-1)K(n\Gamma_{n,\Theta}+V_{n,\Theta})}{2\log(1/1-\beta)\varrho^{\mathrm{AM}}}\right)}{\log(1/a_n^\star)}\right\}$ *and* $K \geq |\bar{\Theta}|\log(|\Theta^\star|/(1-\beta))$*, we have* $\mathbb{P}[\Theta^\star \subseteq \hat{\Theta}_{i,T}^{\mathrm{AM}}] \geq 2\beta - 1$.
- *The optimal distributions that minimize the convergence time* T*, are* $\mathcal{D}_i^\star(\varepsilon) = \mathrm{Lap}\left(\frac{\Delta_{n,\Theta}K|\Theta|}{\varepsilon}\right)$.
- *The beliefs exchanged and the resulting estimates are* ε*-DP with respect to private signals.*

Recall that $\Gamma_{n,\Theta}$ and $l_{n,\Theta}$ that appear in the communication complexity bounds in Theorem 2 are properties of the signal structure defined in Equations (9) and (10)). To minimize communication complexity, it suffices to pick the optimal threshold that makes the terms inside the maximum equal – and, hence, the maximum is minimized – which corresponds to $\varrho_\star^{\mathrm{AM}} = \left(\frac{|\Theta|^2(n-1)K(n\Gamma_{n,\Theta}+V_{n,\Theta})}{2\log(1/(1-\beta))}\right)^{\frac{\log 2}{\log(2/a_n^\star)}}\left(\frac{l_{n,\Theta}}{2n}\right)^{\frac{-\log a_n^\star}{\log(2/a_n^\star)}}$, for AM and $\varrho_\star^{\mathrm{GM}} = \left(\frac{|\Theta|^2(n-1)(n\Gamma_{n,\Theta}+V_{n,\Theta})}{2\alpha\sqrt{K}}\right)^{\frac{\log 2}{\log(2/a_n^\star)}}\left(\frac{l_{n,\Theta}}{2n}\right)^{\frac{-\log a_n^\star}{\log(2/a_n^\star)}}$ for GM.

A Two Threshold Algorithm for Recovering the MLE. In Algorithm 2, we provide a two threshold algorithm to simultaneously control both Type I and Type II error probabilities, with more design flexibility that comes at an increased cost of communication complexity.

The analysis of Theorem 1 shows that in a given round k, the probability that an MLE state $\theta^\star \in \Theta$ ends up positive as $T \to \infty$ is at least $1/|\bar{\Theta}|$, so in expectation at least $K/|\bar{\Theta}|$ rounds will yield a positive belief. On the other hand, we know that for a non-MLE state $\bar{\theta}$, in expectation, at most $(1-1/|\Theta^\star|)K$ trials will come up heads. For brevity, we define $p_1 = 1-1/|\Theta^\star|$ and $p_2 = 1/|\bar{\Theta}|$. Moreover, we define $\boldsymbol{N}_{i,T}^{\mathrm{thres},1}(\hat{\theta})$ (resp. $\boldsymbol{N}_{i,T}^{\mathrm{thres},2}(\hat{\theta})$) as the (average) number of times the belief $\boldsymbol{\nu}_{i,k,t}(\hat{\theta})$ exceeds a threshold $\hat{\tau}^{\mathrm{thres},1}$ (resp. $\hat{\tau}^{\mathrm{thres},2}$) for $k \in [K]$.

Because $\boldsymbol{N}_{i,T}^{\mathrm{thres},1}(\hat{\theta})$ (resp. $\boldsymbol{N}_{i,T}^{\mathrm{thres},2}(\hat{\theta})$) is an average of independent indicator variables for all $\hat{\theta} \in \Theta$, the Chernoff bound indicates that it concentrates around its mean $\mathbb{E}\left[\boldsymbol{N}_{i,T}^{\mathrm{thres},1}(\hat{\theta})\right]$ (resp. $\mathbb{E}\left[\boldsymbol{N}_{i,T}^{\mathrm{thres},2}(\hat{\theta})\right]$). This prompts the development of the following simple algorithm, which resembles boosting algorithms such as AdaBoost (Freund and Schapire, 1997). The algorithm uses two thresholds $\tau^{\mathrm{thres},1}$ and $\tau^{\mathrm{thres},2}$ to estimate the MLE states as those whose beliefs exceed $\hat{\tau}^{\mathrm{thres},1}$ and $\hat{\tau}^{\mathrm{thres},2}$ at least $\tau^{\mathrm{thres},1}$ and $\tau^{\mathrm{thres},2}$ times, leading to output sets $\hat{\Theta}_{i,T}^{\mathrm{thres},1}$ and $\hat{\Theta}_{i,T}^{\mathrm{thres},2}$, respectively.

We first present an asymptotic result (proved in Appendix C.4):

Theorem 3. *For Algorithm 2, with* $p_1 = 1-1/|\Theta^\star|$ *and* $p_2 = 1/|\bar{\Theta}|$*, noise distributions* $\mathcal{D}_i(\hat{\theta};\varepsilon) = \mathcal{D}_i(\varepsilon)$ *that do not depend on the state* $\hat{\theta}$ *and satisfy* ε*-DP, we have that as* $\varrho^{\mathrm{thres},1} \to \infty$ *and* $\varrho^{\mathrm{thres},2} \to \infty$*,*

- *for any* $\pi_1 > 0$*,* $\tau^{\mathrm{thres},1} = (1+\pi_1)p_1$*, and* $K \geq \frac{\log(|\bar{\Theta}|/\alpha)}{2\pi_1^2}$*, we have that* $\lim_{T\to\infty}\mathbb{P}\left[\hat{\Theta}_{i,T}^{\mathrm{thres},1} \subseteq \Theta^\star\right] \geq 1-\alpha$*, for all* $i \in [n]$.

Algorithm 2 Private Distributed MLE (Two Threshold Algorithm)

Inputs: Privacy budget ε, Error probabilities $\alpha, 1-\beta$, Log-belief Thresholds $\varrho^{\text{thres},1}, \varrho^{\text{thres},2} > 0$.

Initialization: Set the number of iterations T and the number of rounds K, and the thresholds $\tau^{\text{thres},2}, \tau^{\text{thres},1}$ as indicated by Theorem 3 (for the asymptotic case) or Theorem 4 (for the non-asymptotic case), and $\hat{\tau}^{\text{thres},1} = \frac{1}{1+e^{\varrho^{\text{thres},1}}}$ (resp. $\hat{\tau}^{\text{thres},2}$). The DP noise distribution for protecting beliefs is $\mathcal{D}_i(\hat{\theta};\varepsilon)$ which can be set optimally according to Theorem 4.

Procedure: The following is repeated in K rounds indexed by $k \in [K]$. In each round k, the agents begin by forming the noisy likelihoods $\boldsymbol{\sigma}_{i,k}(\hat{\theta}) = e^{\boldsymbol{d}_{i,k}(\hat{\theta})}\boldsymbol{\gamma}_i(\hat{\theta})$, where $\boldsymbol{\gamma}_i(\hat{\theta}) = \prod_{j=1}^{n_i} \ell_i(\boldsymbol{s}_i^j|\tilde{\theta})$ and $\boldsymbol{d}_{i,k}(\hat{\theta}) \sim \mathcal{D}_i(\hat{\theta};\varepsilon)$ independently across agents $i \in [n]$ and states $\hat{\theta} \in \Theta$. The agents initialize their beliefs to $\boldsymbol{\nu}_{i,k,0}(\hat{\theta}) = \boldsymbol{\sigma}_{i,k}(\hat{\theta}) / \sum_{\tilde{\theta}\in\Theta} \boldsymbol{\sigma}_{i,k}(\tilde{\theta})$. Over the next T time steps, the agents update their belief after communicating with their neighbors, and according to the following update rule:

$$\boldsymbol{\nu}_{i,k,t}(\hat{\theta}) = \frac{\boldsymbol{\nu}_{i,k,t-1}^{1+a_{ii}}(\hat{\theta}) \prod_{j\in\mathcal{N}_i} \boldsymbol{\nu}_{j,k,t-1}^{a_{ij}}(\hat{\theta})}{\sum_{\tilde{\theta}\in\Theta} \boldsymbol{\nu}_{i,k,t-1}^{1+a_{ii}}(\tilde{\theta}) \prod_{j\in\mathcal{N}_i} \boldsymbol{\nu}_{j,k,t-1}^{a_{ij}}(\tilde{\theta})} \quad \text{for every } \hat{\theta}\in\Theta,\ k\in[K] \text{ and } t\in[T]. \tag{14}$$

After T iterations, each agent aggregates the results of the K rounds as:

$$\boldsymbol{N}_{i,T}^{\text{thres},1}(\hat{\theta}) = \frac{1}{K}\sum_{k\in[K]} \mathbb{1}\left\{\boldsymbol{\nu}_{i,k,T}(\hat{\theta}) > \hat{\tau}^{\text{thres},1}\right\}, \quad \boldsymbol{N}_{i,T}^{\text{thres},2}(\hat{\theta}) = \frac{1}{K}\sum_{k\in[K]} \mathbb{1}\left\{\boldsymbol{\nu}_{i,k,T}(\hat{\theta}) > \hat{\tau}^{\text{thres},2}\right\} \quad \text{for every } \hat{\theta}\in\Theta. \tag{15}$$

Outputs: Return

$$\hat{\Theta}_{i,T}^{\text{thres},1} = \left\{\hat{\theta}\in\Theta : \boldsymbol{N}_{i,T}^{\text{thres},1}(\hat{\theta}) \geq \tau^{\text{thres},1}\right\}, \quad \hat{\Theta}_{i,T}^{\text{thres},2} = \left\{\hat{\theta}\in\Theta : \boldsymbol{N}_{i,T}^{\text{thres},2}(\hat{\theta}) \geq \tau^{\text{thres},2}\right\}.$$

- *for any* $\pi_2 > 0$, $\tau^{\text{thres},2} = (1-\pi_2)p_2$, *and* $K \geq \frac{\log(|\Theta^\star|/(1-\beta))}{2\pi_2^2}$, *we have* $\lim_{T\to\infty} \mathbb{P}\left[\Theta^\star \subseteq \hat{\Theta}_{i,T}^{\text{thres},2}\right] \geq \beta$, *for all* $i \in [n]$.

Similarly to Theorem 1, the estimator $\hat{\Theta}_{i,T}^{1,\text{thres}}$ has a low Type I error, and the estimator $\hat{\Theta}_{i,T}^{2,\text{thres}}$ has a low Type II error. If agents do not know $|\Theta^\star|$, as long as they choose thresholds $\tau^{\text{thres},1'} \geq (1+\pi_1)p_1 = \tau^{\text{thres},1}$ and $\tau^{\text{thres},2'} \leq (1-\pi_2)p_2 = \tau^{\text{thres},2}$, they can obtain the same guarantees. For example, one can choose the following upper and lower bounds to set the thresholds: $\tau^{\text{thres},1'} = (1+\pi_1)(1-1/|\Theta|) \geq (1+\pi_1)p_1$ and $\tau^{\text{thres},2'} = (1-\pi_2)(1/|\Theta|) \leq (1-\pi_2)p_2$. Also, running the algorithm for $K \geq \max\left\{\frac{\log(|\Theta|/\alpha)}{2\pi_1^2}, \frac{\log(|\Theta|/(1-\beta))}{2\pi_2^2}\right\}$ we can guarantee that as $T\to\infty$, we have $\Theta_{i,T}^{\text{thres},2} \subseteq \Theta^\star \subseteq \Theta_{i,T}^{\text{thres},1}$ with probability at least $\beta - \alpha$.

Observe that when $\pi_1 \to \infty$ (which yields $K \to 0$), the result is trivial since the estimator recovered is the empty set which trivially satisfies the error guarantee with probability 1. Similarly, when $\pi_2 \to \infty$ (which again yields $K \to 0$), the guarantee is satisfied with probability one because the entire set Θ is returned.

Because $\boldsymbol{N}_{i,T}^{\text{thres},1}(\theta^\star)$ (resp. $\boldsymbol{N}_{i,T}^{\text{thres},2}(\theta^\star)$) for the MLE states are concentrated around a distribution with a mean that is at least $1/|\bar{\Theta}|$ and, similarly, $\boldsymbol{N}_{i,T}^{\text{thres},1}(\bar{\theta})$ (resp. $\boldsymbol{N}_{i,T}^{\text{thres},2}(\bar{\theta})$) for the non-MLE states

are concentrated around their expectation, which is at most $1 - 1/|\Theta^\star|$, one may rightly question whether it is possible to achieve perfect detection of $\Theta^\star$ assuming that the two modes are "sufficiently" well separated. The answer to this is affirmative as long as $p_2 > p_1$ and is given by the following corollary:

Corollary 1 **(Exact Recovery with a Single Threshold)**. *If the density of MLE states $f^\star$ satisfies:*

$$\begin{cases} 0 \le f^\star \le 1, & \text{if } 1 \le |\Theta| \le 4 \\ 0 \le f^\star < \frac{1}{2} - \frac{1}{2}\sqrt{\frac{|\Theta|-4}{|\Theta|}} \quad \vee \quad \frac{1}{2} + \frac{1}{2}\sqrt{\frac{|\Theta|-4}{|\Theta|}} < f^\star \le 1, & \text{otherwise} \end{cases}, \tag{16}$$

the thresholds are set equal, i.e., $\tau^{\text{thres},2} = \tau^{\text{thres},1}$ and $\hat{\tau}^{\text{thres},1} = \hat{\tau}^{\text{thres},2}$ which corresponds to $\hat{\Theta}_{i,T}^{\text{thres},2} = \hat{\Theta}_{i,T}^{\text{thres},1}$ and $\pi_1 = (1-\pi_2)\frac{p_2}{p_1} - 1$ for some $\pi_2 > 0$, and $K \ge \max\left\{\frac{\log(|\Theta|/\alpha)}{2\pi_1^2}, \frac{\log(|\Theta|/(1-\beta))}{2\pi_2^2}\right\}$, then we have that $\lim_{T\to\infty} \mathbb{P}\left[\hat{\Theta}_{i,T}^{\text{thres},2} = \Theta^\star\right] \ge \beta - \alpha$.

By performing an analysis similar to Theorem 2, we derive a non-asymptotic result (proved in Appendix C.5):

Theorem 4. *Let $q_1, q_2 \in (0,1)$, and $\varrho^{\text{thres},1}, \varrho^{\text{thres},2} > 0$. Then for Algorithm 2 the following hold:*

- *For any $\pi_1 > 0$, $\tau^{\text{thres},1} = (1+\pi_1)q_1$, $T \ge \max\left\{\frac{\log\left(\frac{2\varrho n}{l_{n,\Theta}}\right)}{\log 2}, \frac{\log\left(\frac{|\Theta|^2(n-1)(n\Gamma_{n,\Theta}+V_{n,\Theta})}{2q_1\varrho^{\text{thres},1}}\right)}{\log(1/a_n^\star)}\right\}$, and $K \ge \frac{\log(|\bar{\Theta}|/\alpha)}{2\pi_1^2}$, we have $\mathbb{P}\left[\hat{\Theta}_{i,T}^{\text{thres},1} \subseteq \Theta^\star\right] \ge 1-\alpha$.*
- *For any $\pi_2 > 0$, $\tau^{\text{thres},2} = (1-\pi_2)q_2$, $T \ge \max\left\{\frac{\log\left(\frac{2\varrho n}{l_{n,\Theta}}\right)}{\log 2}, \frac{\log\left(\frac{|\Theta|^2(n-1)(n\Gamma_{n,\Theta}+V_{n,\Theta})}{2(1-q_2)\varrho^{\text{thres},2}}\right)}{\log(1/a_n^\star)}\right\}$, and $K \ge \frac{\log(|\Theta^\star|/(1-\beta))}{2\pi_2^2}$, we have $\mathbb{P}\left[\Theta^\star \subseteq \hat{\Theta}_{i,T}^{\text{thres},2}\right] \ge \beta$.*
- *The optimal distributions that minimize the convergence time T for both estimators are $\mathcal{D}_i^\star(\varepsilon) = \text{Lap}\left(\Delta_{n,\Theta} K|\Theta|/\varepsilon\right)$.*

To minimize the number of iterations, we can pick the thresholds as

$$\varrho_\star^{\text{thres},1} = \left(\frac{|\Theta|^2(n-1)(n\Gamma_{n,\Theta}+V_{n,\Theta})}{2q_1\sqrt{K}}\right)^{\frac{\log 2}{\log(2/a_n^\star)}} \left(\frac{l_{n,\Theta}}{2n}\right)^{\frac{\log(1/a_n^\star)}{\log(2/a_n^\star)}},$$

$$\varrho_\star^{\text{thres},2} = \left(\frac{|\Theta|^2(n-1)(n\Gamma_{n,\Theta}+V_{n,\Theta})}{2(1-q_2)\sqrt{K}}\right)^{\frac{\log 2}{\log(2/a_n^\star)}} \left(\frac{l_{n,\Theta}}{2n}\right)^{\frac{\log(1/a_n^\star)}{\log(2/a_n^\star)}}.$$

In addition, we can show that we can achieve perfect recovery with a single threshold as long as $q_2 > q_1$ and $\pi_1 = (1-\pi_2)\frac{q_2}{q_1} - 1$; similarly to Corollary 1.

Moreover, the above shows that communication complexity can be minimized by setting $1 - q_2$ as close as possible to q_1. However, achieving perfect recovery of MLE comes at a cost to communication complexity, which is polylogarithmic and depends on the inverse of $|1 - q_2 - q_1|$.

To complement our upper bounds for communication complexity, in Appendix G, we give information-theoretic lower bounds on the minimum number of belief exchanges required by any algorithm that achieves the same Type I and Type II error rates as ours. In the next section, we show that the GM estimator can be applied naturally to conduct distributed hypothesis tests at significance level α.

4.2. Differentially Private Distributed Hypothesis Testing

The results we devised above for the MLE have a natural application in the case of hypothesis testing with simple or composite hypotheses. Specifically, the decentralized hypothesis testing algorithm at the significance level α is designed as follows: We run the GM algorithm for T iterations and K rounds and set a threshold ϱ_d to reject the null if the log-belief ratio exceeds it: $(n/2^{T-1})\log(\boldsymbol{\nu}_{i,T}^{\mathrm{GM}}(1)/\boldsymbol{\nu}_{i,T}^{\mathrm{GM}}(0)) > \varrho_d$. For the case of a simple null and a simple alternative, the threshold ϱ_d can be derived from the threshold ϱ_c of the corresponding uniformly most powerful (UMP) centralized log-likelihood ratio test (LRT) at level $\alpha/2$. For sufficiently large T and K we can show that

Proposition 1 **(Simple Hypothesis Testing)**. *Let $\alpha \in (0,1)$ be a significance level. Let ϱ_c be the threshold of the UMP centralized log-likelihood ratio test at level $\alpha/2$, such that $\mathbb{P}[2(\Lambda(1)-\Lambda(0)) \geq \varrho_c|\theta=0] = \alpha/2$. Then by running the GM algorithm (Algorithm 1) with Type I error rate guarantee $\alpha/2$ for T and K given in Theorem 2, and $\varrho_d = \varrho_c - 1$ we get the following: The decentralized test has Type I error at most α, that is, $\mathbb{P}\left[(n/2^{T-1})\log(\boldsymbol{\nu}_{i,T}^{\mathrm{GM}}(1)/\boldsymbol{\nu}_{i,T}^{\mathrm{GM}}(0)) > \varrho_d \middle| \theta=0\right] \leq \alpha$. The resulting hypothesis test is ε-DP.*

The idea behind the result is that, in practice, for sufficiently large T and K, $(n/2^{T-1})\log(\boldsymbol{\nu}_{i,T}^{\mathrm{GM}}(1)/\boldsymbol{\nu}_{i,T}^{\mathrm{GM}}(0)) \sim 2\Lambda(1,0)$ with high probability – i.e., at least $1-\alpha/2$ – due to Chebyshev's inequality. The formal proof of Proposition 1 is in Appendix D.1.

However, in general, when the distribution of the centralized statistic under the null is not available, we cannot choose ϱ_c based on the UMP centralized LRT. In that case, the agents can perform a generalized likelihood ratio test locally and propagate the statistic. This sacrifices the optimality of the UMP test; however, it is more practical in real-world scenarios where the distribution of the sum of the statistics may be unknown under the null hypothesis.

Specifically, the generalized likelihood ratio test can be used to test a composite null hypothesis ($\theta \in \tilde{\Theta}_0$) against a composite alternative ($\theta \in \tilde{\Theta}_1$, $\tilde{\Theta}_0 \cap \tilde{\Theta}_1 = \emptyset$), and requires a log-likelihood function which is twice differentiable defined in a continuous parameter space $\tilde{\Theta} = \tilde{\Theta}_0 \cup \tilde{\Theta}_1$. Specifically, in that case, each agent initializes their belief using $\boldsymbol{\gamma}_i(\tilde{\Theta}_0) = \sup_{\tilde{\theta}_0 \in \tilde{\Theta}_0} \prod_{j=1}^{n_i} \ell_i(\boldsymbol{s}_i^j|\tilde{\theta}_0)$ and $\boldsymbol{\gamma}_i(\tilde{\Theta}) = \sup_{\tilde{\theta} \in \tilde{\Theta}} \prod_{j=1}^{n_i} \ell_i(\boldsymbol{s}_i^j|\tilde{\theta})$, and runs the GM algorithm with the binary state space $\Theta = \{\tilde{\Theta}_0, \tilde{\Theta}\}$. By Wilks' theorem (Casella and Berger, 2002, Theorem 10.3.1), for large sample sizes n_i, each local log-likelihood ratio statistic is asymptotically distributed according to the χ_1^2 distribution, and thus the sum of independent local statistics of tn agents follows a χ_n^2 distribution. We can use the asymptotic distribution of the sum of the local log-likelihood ratio statistics to set the rejection threshold for a distributed level α test to be $\varrho_d = F_{\chi_n^2}^{-1}(1-\alpha/2) - 1$. In contrast, the centralized test has a threshold of $\varrho_c = F_{\chi_1^2}^{-1}(1-\alpha)$. We provide additional details for the extension to composite hypothesis tests in Appendix D.2.

Algorithm 3 Private Distributed Online Learning

Inputs: Privacy budget ε, Error probabilities η.

Initialization: Set the number of iterations T as in Theorem 5. The DP noise distribution for protecting beliefs is $\mathcal{D}_i(\hat{\theta};\varepsilon)$ which can be set optimally according to Theorem 5.

Procedure: Every time $t \in \mathbb{N}_0$, each agent forms the likelihood product of the signals that it has received at that time period: $\boldsymbol{\gamma}_{i,t}(\hat{\theta}) = \prod_{j=1}^{n_{i,t}} \ell_i(\boldsymbol{s}_{i,t}^j|\hat{\theta})$, if $n_{i,t} \geq 1$, and $\gamma_{i,t}(\hat{\theta}) = 1$ if $n_{i,t} = 0$. The agent then draws a noise variable $\boldsymbol{d}_{i,t}(\hat{\theta}) \sim \mathcal{D}_{i,t}(\hat{\theta};\varepsilon)$ for all $i \in [n]$, $\hat{\theta} \in \Theta$ independently and forms $\boldsymbol{\sigma}_{i,t}(\hat{\theta}) = e^{\boldsymbol{d}_{i,t}(\hat{\theta})}\boldsymbol{\gamma}_{i,t}(\hat{\theta})$ for all $\hat{\theta} \in \Theta$. The agent then updates its belief as:

$$\boldsymbol{\nu}_{i,t}(\hat{\theta}) = \frac{\boldsymbol{\sigma}_{i,t}(\hat{\theta})\boldsymbol{\nu}_{i,t-1}^{a_{ii}}(\hat{\theta})\prod_{j\in\mathcal{N}_i}\boldsymbol{\nu}_{j,t-1}^{a_{ij}}(\hat{\theta})}{\sum_{\tilde{\theta}\in\Theta}\boldsymbol{\sigma}_{i,t}(\tilde{\theta})\boldsymbol{\nu}_{i,t-1}^{a_{ii}}(\tilde{\theta})\prod_{j\in\mathcal{N}_i}\boldsymbol{\nu}_{j,t-1}^{a_{ij}}(\tilde{\theta})} \quad \text{for every } \hat{\theta}\in\Theta \text{ and } t\in[T], \tag{17}$$

initialized by: $\boldsymbol{\nu}_{i,0}(\hat{\theta}) = \boldsymbol{\sigma}_{i,0}(\hat{\theta})/\sum_{\tilde{\theta}\in\Theta}\boldsymbol{\sigma}_{i,0}(\tilde{\theta})$.

Outputs: After T iterations, return $\hat{\theta}_{i,T}^{\mathrm{OL}} = \arg\max_{\hat{\theta}\in\Theta}\boldsymbol{\nu}_{i,T}(\hat{\theta})$.

4.3. Differentially Private Distributed Online Learning from Intermittent Streams

Algorithm 3 facilitates differentially private distributed learning in an online setting, where agents receive $n_{i,t}$ signals at each time step t, calculate the likelihood $\boldsymbol{\gamma}_{i,t}(\hat{\theta}) = \prod_{j=1}^{n_{i,t}} \ell_i(\boldsymbol{s}_{i,t}^j|\hat{\theta})$ and its noisy version $\boldsymbol{\sigma}_{i,t}(\hat{\theta}) = e^{\boldsymbol{d}_{i,t}(\hat{\theta})}\boldsymbol{\gamma}_{i,t}(\hat{\theta})$ to preserve DP, and then aggregate it with their self- and neighboring beliefs from iteration $t-1$. Unlike the MLE setup, for online learning, we do not need to repeat the algorithm in multiple rounds because DP noise cancels out as $T \to \infty$.

The asymptotic dynamics of beliefs under Algorithm 3 becomes clear when you notice that the log-belief ratio between θ° and any other state $\bar{\theta} \neq \theta^\circ$ converges to $(-1/n)\sum_{i=1}^n \xi_i D_{KL}\left(\ell_i(\cdot|\bar{\theta})|\ell_i(\cdot|\theta^\circ)\right)$ as $t \to \infty$ and the effect of DP noise disappears as a consequence of the Césaro mean and the weak law of large numbers. Therefore, agents learn the true states θ° exponentially fast, asymptotically (one can show that the asymptotic exponential rate is $-l_{n,\Theta}/n$, where $l_{n,\Theta}$ is given by Equation (10)). The next theorem characterizes the nonasymptotic behavior of Algorithm 3 (proved in Appendix E):

Theorem 5. *Given*

$$T = \mathcal{O}\left(\frac{n|\Theta|(Q_{n,\Theta} + \Delta_{n,\Theta}|\Theta|/\varepsilon)\left(\max_{i\in[n]}\xi_i + \sqrt{\frac{\Xi_n}{\eta}}\right)}{l_{n,\Theta}\eta(1-b_n^\star)}\right)$$

iterations, where $l_{n,\Theta}$ and $Q_{n,\Theta}$ are properties of the signal structure defined in Equations (10) *and* (11)*, running Algorithm 3 yields $\mathbb{P}\left[\theta^\circ = \hat{\theta}_{i,T}^{\mathrm{OL}}\right] \geq 1-\eta$ as long as $\sum_{i=1}^n \xi_i D_{KL}\left(\ell_i(\cdot|\bar{\theta})|\ell_i(\cdot|\theta^\circ)\right) > 0$ for all $\bar{\theta} \neq \theta^\circ$. Subsequently, the noise distributions that optimize runtime are $\mathcal{D}_{i,t}^\star = \mathrm{Lap}\left(\Delta_{n,\Theta}|\Theta|/\varepsilon\right)$. The resulting estimates are ε-DP with respect to the private signals.*

Note that similarly to distributed MLE algorithms, here we also need to scale the privacy budget by $|\Theta|$ to account for the fact that the algorithm accesses the private signals and adds noise for each of the $|\Theta|$ states, providing an attacker with as many observations of the private signals.

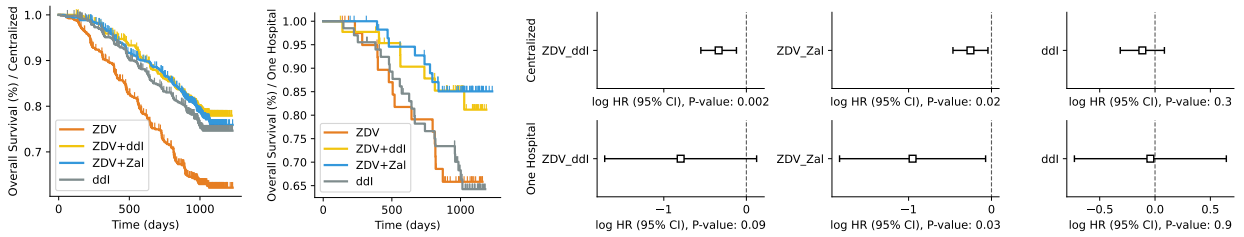

(a) Centralized (b) One hospital (c) Log Hazard Ratios

Figure 4 **Left:** Centralized curves. **Middle:** Curves for one hospital. Data is split evenly among $n = 5$ hospitals. **Right:** Log Hazard Ratios for each of the treatments for centralized and for one hospital fitted on the data. 95% confidence intervals are reported. The data is split uniformly across centers.

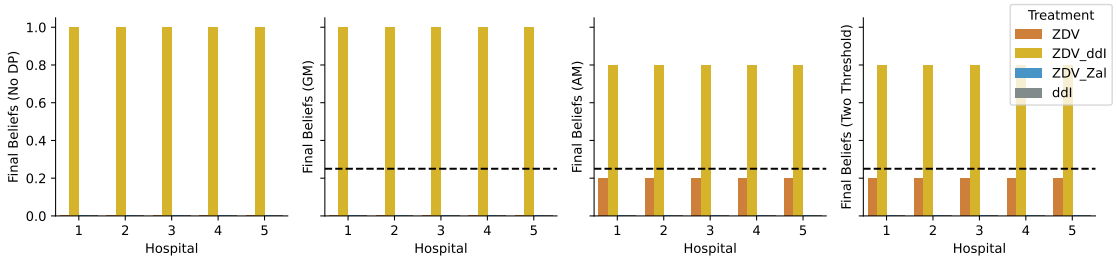

(a) Final Beliefs

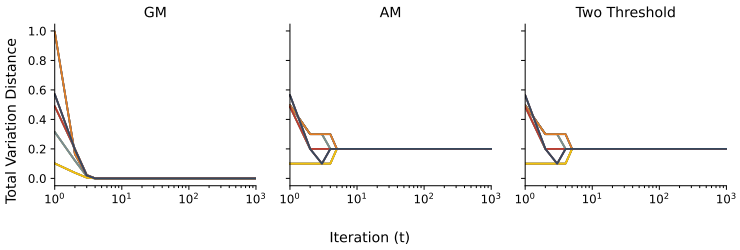

(b) Total Variation Distance

Figure 5 **Top:** Resulting beliefs (at terminal time T) for distributed MLE for AM, GM, and the Two-Threshold algorithm (recovery is performed with one threshold). **Bottom:** Total variation distance between the results of each algorithm and the non-DP baseline. The privacy budget is set to be $\varepsilon = 1$, and the errors to be $\alpha = 1 - \beta = 0.05$. The thresholds are set to 0.25 (in the belief space). The network corresponds to a network of $n = 5$ fully connected centers. All algorithms recover the best treatment (ZDV+ddI; see also Figure 4).

5. Simulation Study

For the simulation study, we focus on the AIDS clinical trial dataset. Due to space constraints, several additional results (including simulations with a different dataset of cancer patients) are referred to here and deferred to Appendix H.

In the ACTG dataset, the control is zidovudine (ZDV), and the three possible treatments are didanosine (ddI), zidovudine plus didanosine (ZDV + ddI), and zidovudine plus zalcitabine (ZDV +

Zal). The survival curves and log hazard ratios from the fitted proportional hazards model are shown in Figure 4.

Hypothesis testing with a single treatment. Initially, we are interested in whether a single treatment improves patient survival. Recall from Section 2.1, that the treatment variable is modeled as a covariate $\boldsymbol{x}_{ij} \in \{0,1\}$. We are interested in the following hypothesis test with a simple null and a composite alternative where we have used ddI as a treatment and ZDV as the control/standard care (similar results for simple hypothesis testing $\theta = 0$ vs $\theta = -\log 2$ are provided in Appendix H.1):

$\tilde{\Theta}_0\ (\theta = 0)$: *No difference between patient survivals under the standard care and the new treatment.*

$\tilde{\Theta}_1\ (\theta < 0)$: *The new treatment improves patient survival compared to the standard care.*

We consider a network of $n = 5$ fully connected hospitals and a significance level of $\alpha = 0.05$, which is usually used in cases of clinical trials, run the algorithm presented in Appendix D and compare against the centralized hypothesis testing algorithm. We evenly split the data between hospitals and between the control and treatment groups (that is, hospitals with the same amount of control or treated patients). In Appendix F we calculate the global sensitivity $\Delta_{n,\Theta}$ of the proportional hazards model and show that it is $\Delta_{n,\Theta} = 2B_\theta$ where B_θ is an upper bound on the parameter value. We use a privacy budget of $\varepsilon = 1$. Generalized likelihood ratio statistics have been calculated with the open-source Python package `lifelines` (Davidson-Pilon, 2019), which offers methods to fit the proportional hazards model. To improve the stability of estimates and controls for the high correlation between covariates, we have used an L_2 regularization term $\lambda_{\text{reg}}|\theta|^2$ with $\lambda_{\text{reg}} = 0.05$.

The centralized test yields a p-value of $P < 10^{-4}$. Similarly, the distributed hypothesis test produces a p-value of $P < 10^{-5}$, showing that treatment (ddI) positively affects survival.

Best treatment identification among multiple treatments. Next, we leverage our framework to identify the best of the three treatments. To achieve that in practice with a numerically stable algorithm, we propagate the generalized likelihood ratios over each of the treatments and then convert them to beliefs. We use the same topology as above (a complete network of $n = 5$ centers), set $\varepsilon = 1$, and the errors to be $\alpha = 1 - \beta = 0.05$. Figure 5 shows the resulting beliefs as well as the total variation distance between each of the inference algorithms and the non-DP baseline. All algorithms are able to identify the best treatment that corresponds to ZDV + ddI and is in agreement with the ground truth (cf. Figure 4).

Runtime study. In Appendix H.2, we characterize the computational complexity of our algorithms, running in polynomial time and space. In Appendix H.3, we perform a runtime study on our proposed algorithms and show that our algorithm can run between $\sim 10^{-2}$ and 10 seconds for values of n ranging from $n = 3$ to $n = 96$, which is significantly faster (up to $1000\times$) than existing methods that rely on homomorphic encryption (Froelicher et al., 2021). Fully homomorphic encryption and multiparty computation are cryptographic techniques that allow computation on encrypted data

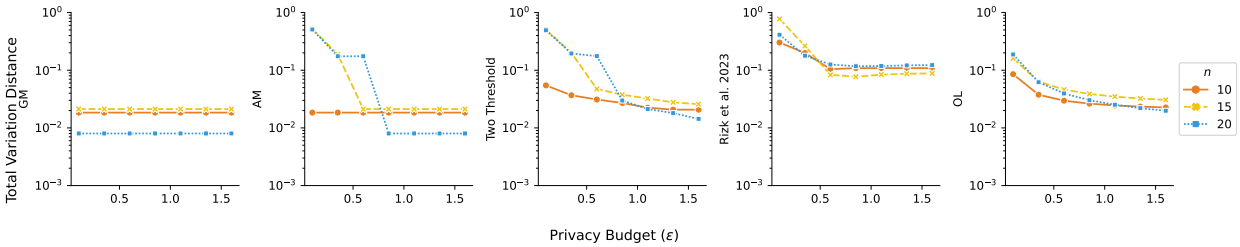

Figure 6 Average total variation distance between the algorithm outputs and the ground truth for the MLE algorithms (AM/GM/Two-threshold), the OL algorithm, and the first-order method introduced in Rizk et al. (2023) for a range of values for the privacy budget ε and $n \in \{10, 15, 20\}$ centers. Our MLE algorithms exhibit a smaller average total variation distance compared to the algorithm of Rizk et al. (2023).

(by single or multiple severs, respectively) without the need to access decrypted data. The strong cryptographic guarantees of these methods come at a high cost to computational efficiency, and their implementation requires information technology (IT) infrastructure and large-scale, high-performance computational resources across distributed sites, which are beyond the technical capabilities of many organizations.

Comparison with first-order methods. First-order distributed optimization methods with differential privacy guarantees offer an alternative approach Rizk et al. (2023), which, while more general, can be inefficient for distributed maximum likelihood estimation. Figure 6 shows that our algorithms yield answers closer to ground truth compared to the first-order method of Rizk et al. (2023) with up to 100 times smaller errors (details in Appendix H.4).

Additional datasets and experiments. In Appendix H.5, we provide experiments with data from clinical trials in patients with advanced cancer, where the task is to determine whether certain biometric indices affect patient survival.

6. Discussion

In theory, introducing privacy in the distributed MLE task incurs a $\text{polylog}\,(1/\varepsilon, 1/\eta)$ cost, for the AM/GM algorithms and the two-threshold algorithms, compared to the non-private benchmark with respect to ε and η, and a $\text{poly}\,(|\Theta|, \log \Delta_{n,\Theta})$ cost compared to the non-private benchmark with respect to $|\Theta|$, and a $\text{polylog}\,(n)$ overhead with respect to n.

Moreover, in the online learning regime, we end up having a $\text{poly}\,(1/\varepsilon, 1/\eta)$ dependency compared to the non-private benchmark with respect to ε and η, and a $\text{poly}\,(|\Theta|, \Delta_{n,\Theta})$ dependency with respect to $|\Theta|$. This shows that achieving the same privacy and accuracy levels in the online case requires many more communication rounds than in the MLE case.

Societal impacts. Our article touches on topics at the intersections of group decision making and privacy. Although our work is primarily theoretical and does not pose specific ethical challenges

to the best of our knowledge, its societal and policy impact should be highlighted. Our work has applications for distributed hypothesis testing problems, such as, for example, testing new treatments in which multiple entities want to estimate the actual state of nature based on observations from their neighbors without exposing their information, and provide protocols for information fusion depending on the application. Work on such sensitive applications may require the development or adaptation of privacy regulations. Policymakers may need to consider balancing the benefits of collaborative decision making with protecting individual privacy and intellectual property, especially given that our algorithms provide the correct answer with a high probability (but are still susceptible to negligible errors). Our paper's approach allows entities to benefit from each other's private information without compromising data security. As a result, organizations can improve the quality of their services through better decision making without violating privacy obligations, positively impacting efficiency and cost savings.

Concluding remarks. In summary, in this paper, we study distributed estimation and learning in network environments where agents face privacy risks with respect to their own data and the beliefs expressed by their neighbors. We propose two algorithms for distributed MLE for which we give guarantees to control Type I and Type II errors. Our update rules are based on the decision-analysis literature and possess certain optimality guarantees. The agents add noise to their own signals and neighboring beliefs (depending on the nature of the protections) and exchange their beliefs with each other. To control the estimation errors, we devise upper bounds on the number of iterations and the number of times the algorithms should be applied, which can be thought of as their communication complexity. Introducing privacy-preserving computations incurs trade-offs in communication complexity and shows that the mechanism that optimizes the privacy overhead is the Laplace mechanism with appropriately chosen parameters. Finally, we test our algorithms and validate our theory in experiments with real-world data from multicenter clinical trials.

Data and Code Availability Statement

The data used to support the findings of the paper can be downloaded from the following sources:

1. The AIDS Clinical Trials data set was obtained from Kaggle: `https://www.kaggle.com/datasets/tanshihjen/aids-clinical-trials` (accessed February 14, 2025). Original information about the clinical trial can be found at `https://clinicaltrials.gov/study/NCT00000625` (accessed February 14, 2025).
2. The data set for the cancer clinical trial was obtained from Samstein et al. (2019).

The codes for reproducing the simulation results and figures are open source and can be accessed at `https://github.com/papachristoumarios/dp-social-learning`.

Acknowledgments

M.P. was partially supported by a LinkedIn Ph.D. Fellowship, an Onassis Fellowship (ID: F ZT 056-1/2023-2024), and grants from the A.G. Leventis Foundation and the Gerondelis Foundation. M.A.R. was partially supported by NSF SaTC-2318844. The authors would like to thank the seminar participants at Rutgers Business School, Jalaj Upadhyay, Saeed Sharifi-Malvajerdi, Jon Kleinberg, Kate Donahue, Richa Rastogi, the attendants of the Information Systems Workshop at Arizona State University, Sue Brown, Maytal Saar-Tschechansky, and Vasilis Charisopoulos for their valuable discussions and feedback.

References

Abbas, Ali E (2009). A Kullback-Leibler view of linear and log-linear pools. *Decision Analysis* 6(1): 25–37.

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asu Ozdaglar (2022). Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Microeconomics* 14(4): 218–256.

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar (2017). Privacy-constrained network formation. *Games and Economic Behavior* 105: 255–275.

Acquisti, Alessandro, Laura Brandimarte, and George Loewenstein (2015). Privacy and human behavior in the age of information. *Science* 347(6221): 509–514.

Adjerid, Idris, Eyal Peer, and Alessandro Acquisti (2018). Beyond the Privacy Paradox. *MIS quarterly* 42(2): 465–488.

Apple Differential Privacy Team (2017). Learning with Privacy at Scale. `https://machinelearning.apple.com/research/learning-with-privacy-at-scale`. Accessed: 2023-05-18.

ARPA-H (2024). ARPA-H Launches Groundbreaking Funding Opportunity to Improve Clinical Trials. `https://arpa-h.gov/news-and-events/arpa-h-launches-groundbreaking-funding-opportunity-improve-clinical-trials`. Accessed: 2024-12-13.

Aumann, R. J. (1976). Agreeing to disagree. *The annals of statistics*: 1236–1239.

Awan, Jordan and Aleksandra Slavković (2018). Differentially private uniformly most powerful tests for binomial data. *Advances in Neural Information Processing Systems* 31.

Banerjee, Abhijit, Emily Breza, Arun G Chandrasekhar, and Markus Mobius (2021). Naive learning with uninformed agents. *American Economic Review* 111(11): 3540–3574.

Barbaro, Michael, Tom Zeller, and Saul Hansell (2006). A face is exposed for AOL searcher no. 4417749. *New York Times* 9(2008): 8.

Bélanger, France and Robert E Crossler (2011). Privacy in the digital age: a review of information privacy research in information systems. *MIS quarterly*: 1017–1041.

Bélanger, France and Tabitha L James (2020). A theory of multilevel information privacy management for the digital era. *Information systems research* 31(2): 510–536.

Benthall, Sebastian and Rachel Cummings (2022). Integrating differential privacy and contextual integrity. *2022 USENIX Conference on Privacy Engineering Practice and Respect*. Santa Clara, CA, USA.

Blum, Avrim, Jamie Morgenstern, Ankit Sharma, and Adam Smith (2015). Privacy-preserving public information for sequential games. *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science*: 173–180.

Borkar, V. and P.P. Varaiya (June 1982). Asymptotic agreement in distributed estimation. *IEEE Transactions on Automatic Control* 27(3): 650–655.

Boyd, Stephen, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah (2006). Randomized gossip algorithms. *IEEE Transactions on Information Theory* 52(6): 2508–2530.

Bureau, US Census (2021). 2020 census results.

Cai, Bryan, Constantinos Daskalakis, and Gautam Kamath (2017). Priv'it: Private and sample efficient identity testing. *International Conference on Machine Learning*. PMLR: 635–644.

Campbell, Thomas B et al. (2012). Efficacy and safety of three antiretroviral regimens for initial treatment of HIV-1: a randomized clinical trial in diverse multinational settings. *PLoS One*.

Canonne, Clément L et al. (2019). The structure of optimal private tests for simple hypotheses. *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*: 310–321.

Casella, George and Roger L Berger (2002). Statistical inference. Vol. 2. Duxbury Pacific Grove, CA.

Clemen, Robert T and Robert L Winkler (1999). Combining probability distributions from experts in risk analysis. *Risk analysis* 19: 187–203.

Conradt, Larissa, Christian List, and Timothy J Roper (2013). Swarm intelligence: When uncertainty meets conflict. *The American Naturalist* 182(5): 592–610.

Cortes, J., S. Martinez, and F. Bullo (June 2005). Analysis and design tools for distributed motion coordination. *Proceedings of the American Control Conference*. Portland, OR: 1680–1685.

Cox, David R (1972). Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)* 34(2): 187–202.

Dahleh, M, Alireza Tahbaz-Salehi, John N Tsitsiklis, and Spyros I Zoumpoulis (2012). On global games in social networks of information exchange. Tech. rep. Working paper.

Davidson-Pilon, Cameron (2019). lifelines: survival analysis in Python. *Journal of Open Source Software* 4(40): 1317.

DeGroot, M. H. (1974). Reaching a consensus. *Journal of American Statistical Association* 69: 118–121.

DeMarzo, P. M., D. Vayanos, and J. Zwiebel (2003). Persuasion bias, social influence, and unidimensional opinions. *The Quarterly Journal of Economics* 118: 909–968.

Diana, Emily et al. (2020). Differentially private call auctions and market impact. *Proceedings of the 21st ACM Conference on Economics and Computation*: 541–583.

Dietrich, Franz, Christian List, and Richard Bradley (2016). Belief revision generalized: A joint characterization of Bayes' and Jeffrey's rules. *Journal of Economic Theory* 162: 352–371.

Dwork, Cynthia, Nitin Kohli, and Deirdre Mulligan (2019). Differential privacy in practice: Expose your epsilons! *Journal of Privacy and Confidentiality* 9(2).

Dwork, Cynthia and Aaron Roth (2014). The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science* 9(3–4): 211–407.

Erlingsson, U (2014). Learning statistics with privacy, aided by the flip of a coin. *Google Security Blog, October*.

Ficek, Joseph et al. (2021). Differential privacy in health research: A scoping review. *Journal of the American Medical Informatics Association* 28(10): 2269–2276.

Freund, Yoav and Robert E Schapire (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences* 55(1): 119–139.

Friedkin, Noah E and Eugene C Johnsen (1990). Social influence and opinions. *Journal of Mathematical Sociology* 15(3-4): 193–206.

Froelicher, David et al. (2021). Truly privacy-preserving federated analytics for precision medicine with multiparty homomorphic encryption. *Nature communications* 12(1): 5910.

Garfinkel, Simson (2022). Differential privacy and the 2020 US census. Tech. rep. MIT Schwarzman College of Computing.

Garthwaite, Paul H, Joseph B Kadane, and Anthony O'Hagan (2005). Statistical methods for eliciting probability distributions. *Journal of the American Statistical Association* 100(470): 680–701.

Geanakoplos, J. D. and H. M. Polemarchakis (1982). We can't disagree forever. *Journal of Economic Theory* 28(1): 192–200.

Genest, Christian, Samaradasa Weerahandi, and James V Zidek (1984). Aggregating opinions through logarithmic pooling. *Theory and Decision* 17(1): 61–70.

Genest, Christian and James V Zidek (1986). Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*: 114–135.

Geva, Ravit et al. (2023). Collaborative privacy-preserving analysis of oncological data using multiparty homomorphic encryption. *Proceedings of the National Academy of Sciences* 120(33): e2304415120.

Gilardoni, G. L. and M. K. Clayton (1993). On reaching a consensus using DeGroot's iterative pooling. *The Annals of Statistics*: 391–401.

Golub, B. and M. O. Jackson (Feb. 2010). Naïve Learning in Social Networks and the Wisdom of Crowds. *American Economic Journal: Microeconomics* 2(1): 112–149.

Golub, Benjamin and Matthew O. Jackson (2012). How Homophily Affects the Speed of Learning and Best-Response Dynamics. *The Quarterly Journal of Economics* 127(3): 1287–1338.

Hammer, Scott M et al. (1996). A trial comparing nucleoside monotherapy with combination therapy in HIV-infected adults with CD4 cell counts from 200 to 500 per cubic millimeter. *New England Journal of Medicine* 335(15): 1081–1090.

Hatano, Y. and M. Mesbahi (2005). Agreement over random networks. *IEEE Transactions on Automatic Control* 50(11): 1867–1872.

Hegselmann, R. and U. Krause (2002). Opinion dynamics and bounded confidence models, analysis and simulation. *Journal of Societies and Social Simulation* 5(3).

Jadbabaie, A., J. Lin, and A. S. Morse (2003). Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control* 48(6): 988–1001.

Jadbabaie, A., P. Molavi, A. Sandroni, and A. Tahbaz-Salehi (2012). Non-Bayesian social learning. *Games and Economic Behavior* 76(1): 210–225.

Kang, Justin, Ramtin Pedarsani, and Kannan Ramchandran (2023). The Fair Value of Data Under Heterogeneous Privacy Constraints. *arXiv preprint arXiv:2301.13336*.

Kaplan, Edward L and Paul Meier (1958). Nonparametric estimation from incomplete observations. *Journal of the American statistical association* 53(282): 457–481.

Kayaalp, Mert, Yunus Inan, Emre Telatar, and Ali H Sayed (2023). On the arithmetic and geometric fusion of beliefs for distributed inference. *IEEE Transactions on Automatic Control* 69(4): 2265–2280.

Kazan, Zeki, Kaiyan Shi, Adam Groce, and Andrew P Bray (2023). The test of tests: A framework for differentially private hypothesis testing. *International Conference on Machine Learning*. PMLR: 16131–16151.

Kearns, Michael, Mallesh Pai, Aaron Roth, and Jonathan Ullman (2014). Mechanism design in large games: Incentives and privacy. *Proceedings of the 5th conference on Innovations in theoretical computer science*: 403–410.

Koufogiannis, Fragkiskos, Shuo Han, and George J Pappas (2015). Optimality of the laplace mechanism in differential privacy. *arXiv preprint arXiv:1504.00065*.

Kumar, Ravi, Jasmine Novak, Bo Pang, and Andrew Tomkins (2007). On anonymizing query logs via token-based hashing. *Proceedings of the 16th international conference on World Wide Web*: 629–638.

Lalitha, Anusha, Tara Javidi, and Anand D Sarwate (2018). Social learning and distributed hypothesis testing. *IEEE Transactions on Information Theory* 64(9): 6161–6179.

Lehrer, Keith and Carl Wagner (1981). Rational Consensus in Science and Society, A Philosophical and Mathematical Study. Dordrecht, Holland: D. Reidel Publishing Company.

Li, Chencheng et al. (2018). Differentially private distributed online learning. *IEEE Transactions on Knowledge and Data Engineering* 30(8): 1440–1453.

Liell-Cock, Jack, Ian R Manchester, and Guodong Shi (2020). Preserving privacy of the influence structure in Friedkin-Johnsen systems. *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE: 6254–6259.

List, Christian and Clemens Puppe (2009). Judgment aggregation: A survey. *Handbook of Rational and Social Choice*.

Marden, Jason R and Jeff S Shamma (2012). Revisiting log-linear learning: Asynchrony, completeness and payoff-based implementation. *Games and Economic Behavior* 75(2): 788–808.

McLaughlin, Samuel, Vikram Krishnamurthy, and Subhash Challa (2003). Managing data incest in a distributed sensor network. *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'03)*. Vol. 5. IEEE: V–269.

Mesbahi, M. and M. Egerstedt (2010). Graph Theoretic Methods in Multiagent Networks. Princeton University Press.

Molavi, Pooya, Alireza Tahbaz-Salehi, and Ali Jadbabaie (2018). A theory of non-Bayesian social learning. *Econometrica* 86(2): 445–490.

Moreau, L. (2005). Stability of Multiagent Systems With Time-Dependent Communication Links. *IEEE Transactions on Automatic Control* 50(2): 169–182.

Munjal, Kundan and Rekha Bhatia (2023). A systematic review of homomorphic encryption and its contributions in healthcare industry. *Complex & Intelligent Systems* 9(4): 3759–3786.

National Institute of Allergy and Infectious Diseases (1995). A Randomized, Double-Blind Phase II/III Trial of Monotherapy vs. Combination Therapy With Nucleoside Analogs in HIV-Infected Persons With CD4 Cells of 200-500/mm3. `https://clinicaltrials.gov/study/NCT00000625`. Accessed: 2024-12-13.

Nissenbaum, Helen (2009). Privacy in context. *Privacy in Context*. Stanford University Press.

Nissim, Kobbi, Sofya Raskhodnikova, and Adam Smith (2007). Smooth sensitivity and sampling in private data analysis. *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*: 75–84.

Olshevsky, Alex (2017). Linear time average consensus and distributed optimization on fixed graphs. *SIAM Journal on Control and Optimization* 55(6): 3990–4014.

Papachristou, Marios and M Amin Rahimian (2024). Differentially Private Distributed Estimation and Learning. *IISE Transactions* (just-accepted): 1–26.

Qin, Tiancheng, S Rasoul Etesami, and Cesár A Uribe (2022). Decentralized federated learning for over-parameterized models. *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE: 5200–5205.

Rabbat, Michael G, Robert D Nowak, and James A Bucklew (2005). Generalized consensus computation in networked systems with erasure links. *Signal Processing Advances in Wireless Communications, 2005 IEEE 6th Workshop on*. IEEE: 1088–1092.

Rahimian, M Amin and Ali Jadbabaie (2016a). Bayesian learning without recall. *IEEE Transactions on Signal and Information Processing over Networks* 3(3): 592–606.

— (2016b). Distributed estimation and learning over heterogeneous networks. *Communication, Control, and Computing (Allerton), 2016 54th Annual Allerton Conference on*. IEEE: 1314–1321.

Rahimian, M Amin, Fang-Yi Yu, and Carlos Hurtado (2023a). Differentially Private Network Data Collection for Influence Maximization. *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*: 2795–2797.

— (2023b). Seeding with Differentially Private Network Information. *arXiv preprint arXiv:2305.16590*.

Rahimian, M. A. and A. Jadbabaie (2015). Learning without recall: A case for log-linear learning. *IFAC-PapersOnLine* 48(22): 46–51.

Rahimian, M. A., S. Shahrampour, and A. Jadbabaie (2015). Learning without Recall by Random Walks on Directed Graphs. *IEEE Conference on Decision and Control (CDC)*.

Rehm, Heidi L et al. (2021). GA4GH: International policies and standards for data sharing across genomic research and healthcare. *Cell genomics* 1(2).

Ren, W. and RW Beard (2005). Consensus seeking in multiagent systems under dynamically changing interaction topologies. *IEEE Transactions on Automatic Control* 50(5): 655–661.

Rezaei, Aria, Jie Gao, and Anand D Sarwate (2021). Influencers and the Giant Component: The Fundamental Hardness in Privacy Protection for Socially Contagious Attributes. *Proceedings of the 2021 SIAM International Conference on Data Mining (SDM)*. SIAM: 217–225.

Rizk, Elsa, Stefan Vlaski, and Ali H Sayed (2023). Enforcing privacy in distributed learning with performance guarantees. *IEEE Transactions on Signal Processing* 71: 3385–3398.

Rufo, M. J., J. Martin, C. J. Pérez, et al. (2012). Log-Linear Pool to Combine Prior Distributions: A Suggestion for a Calibration-Based Approach. *Bayesian Analysis* 7(2): 411–438.

Samstein, Robert M et al. (2019). Tumor mutational load predicts survival after immunotherapy across multiple cancer types. *Nature genetics* 51(2): 202–206.

Shahrampour, Shahin, Alexander Rakhlin, and Ali Jadbabaie (2015). Distributed detection: Finite-time analysis and impact of network topology. *IEEE Transactions on Automatic Control* 61(11): 3256–3268.

Slivkins, Aleksandrs et al. (2019). Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning* 12(1-2): 1–286.

Stoica, Ana-Andreea and Christos Papadimitriou (2021). Strategic clustering. *Learning in the presence of strategic behaviour Workshop (StratML), in Neural Information Processing Systems*.

Sweeney, Latanya (1997). Weaving technology and policy together to maintain confidentiality. *The Journal of Law, Medicine & Ethics* 25(2-3): 98–110.

— (2015). Only you, your doctor, and many others may know. *Technology Science* 2015092903(9): 29.

Tanner, H., A. Jadbabaie, and G. Pappas (2007). Flocking in Fixed and Switching Networks. *IEEE Transactions on Automatic Control* 52(5): 863–868.

Tsitsiklis, J. N. and M. Athans (1984). Convergence and asymptotic agreement in distributed decision problems. *Automatic Control, IEEE Transactions on* 29(1): 42–50.

Wan, Zhiyu et al. (2022). Sociotechnical safeguards for genomic data privacy. *Nature Reviews Genetics* 23(7): 429–445.

Warner, Stanley L (1965). Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association* 60(309): 63–69.

Wu, Bohan and César A Uribe (2023). Frequentist Guarantees of Distributed (Non)-Bayesian Inference. *arXiv preprint arXiv:2311.08214*.

Xiao, L., S. Boyd, and S. Lall (2005). A scheme for robust distributed sensor fusion based on average consensus. *Fourth International Symposium on Information Processing in Sensor Networks (IPSN)*: 63–70.

Youn, Yeojoon, Zihao Hu, Juba Ziani, and Jacob Abernethy (2023). Randomized quantization is all you need for differential privacy in federated learning. *arXiv preprint arXiv:2306.11913*.

Zhao, Huimin and Sudha Ram (2007). Combining schema and instance information for integrating heterogeneous data sources. *Data & Knowledge Engineering* 61(2): 281–303.

Ziani, Juba (2022). How Differential Privacy Impacts Data Elicitation. *ACM SIGecom Exchanges* 20(2): 75–81.

Appendix

Appendix A: Additional Related Work

In this section, we review the following collections of literature on privacy and group decision making: *(i)* differential privacy, *(ii)* works at the intersection of group decision making and social learning, *(iii)* privacy protections in social network contexts, *(iv)* literature in distributed learning and estimation, and *(v)* existing advances in the literature of privacy in healthcare contexts.

Differential privacy. Data privacy has evolved to incorporate various concepts, including K-anonymity. In modern definition, a mechanism is considered differentially private if it assigns similar records to the same value with equal likelihood. This definition has a significant implication. It ensures that the outcome of a statistical analysis remains consistent regardless of an individual's participation in the social learning process. Numerous existing mechanisms, such as the randomized response (Warner, 1965), the Laplace mechanism, and the Gaussian mechanism, can be demonstrated to adhere to differential privacy principles. The randomized response, for example, introduces random perturbations in binary responses, allowing the retrieval of population means while allowing individual respondents to maintain plausible deniability. It serves as an example of adding noise to the data.

When it comes to handling donated data for purposes such as providing recommendations to public policy agencies or submitting online product reviews on e-commerce platforms, it is crucial to protect the privacy of data donors. This is because common anonymization techniques used in such scenarios are susceptible to different types of attacks, including identification risks (Barbaro et al., 2006), linkage and cross-referencing vulnerabilities (Sweeney, 1997; Sweeney, 2015), as well as statistical difference and reidentification attacks (Kumar et al., 2007). Therefore, protecting the privacy of data donors becomes an essential aspect of statistical disclosure control in these contexts with important trade-offs between privacy and utility (Ziani, 2022). Our work contributes to the literature on DP by utilizing DP techniques to enable group decision making in distributed environments.

Group decision making and social learning. The problem of aggregating the opinions and observations of different individuals into a coherent collective option arises naturally in jury deliberations, expert committees, medical diagnoses, or board meetings where several stakeholders must agree on a factual basis despite heterogeneity and uncertainty of their opinions and private information. A common way to resolve disagreements in such contexts is to engage in repeated peer interactions. Subsequently, a long line of literature goes back to seminal works of (Aumann, 1976; Geanakoplos and Polemarchakis, 1982; DeGroot, 1974) that investigate the formation and evolution of Bayesian and non-Bayesian beliefs towards consensus agreement. This problem has close parallels in distributed estimation (Borkar and Varaiya, 1982; Tsitsiklis and Athans, 1984), data fusion (McLaughlin et al.,

2003), as well as consensus and coordination in distributed control (Jadbabaie et al., 2003; Mesbahi and Egerstedt, 2010). The formation, evolution and aggregation of beliefs in distributed environments have attracted a lot of interest in statistics (Genest et al., 1984; Genest and Zidek, 1986; Garthwaite et al., 2005; Wu and Uribe, 2023; Qin et al., 2022), engineering (Kayaalp et al., 2023; Moreau, 2005; Hatano and Mesbahi, 2005; Boyd et al., 2006; Xiao et al., 2005; Ren and Beard, 2005; Cortes et al., 2005; Rabbat et al., 2005; Tanner et al., 2007; Borkar and Varaiya, 1982; Tsitsiklis and Athans, 1984; Shahrampour et al., 2015), philosophy (Lehrer and Wagner, 1981; List and Puppe, 2009; Dietrich et al., 2016; Conradt et al., 2013), sociology (Hegselmann and Krause, 2002; Friedkin and Johnsen, 1990), machine learning (Stoica and Papadimitriou, 2021), and economics (Golub and Jackson, 2010; DeMarzo et al., 2003; Golub and Jackson, 2012; Jadbabaie et al., 2012; Banerjee et al., 2021; Diana et al., 2020; Dahleh et al., 2012), among others.

Of particular interest to us in this paper is the growing literature on non-Bayesian information aggregation and opinion pooling, going back to the classical work of DeGroot (DeGroot, 1974), who proposes linear opinion pools by averaging the latest beliefs of each agent with those of their neighbors. Several asymptotic properties and extensions of this model, including streams of private signals over time, are studied in the literature (Jadbabaie et al., 2012; Golub and Jackson, 2010; DeMarzo et al., 2003). To combine beliefs on a finite set of alternatives, we use geometric averaging and logarithmic opinion pools, which also have a long history in Bayesian analysis and behavioral decision models (Gilardoni and Clayton, 1993; Rufo et al., 2012) and can be justified under specific axioms (Molavi et al., 2018) or derived as a no-recall behavioral update rule (Rahimian and Jadbabaie, 2015; Rahimian et al., 2015; Rahimian and Jadbabaie, 2016a). In this paper, we ask how one should limit the information requirement of a non-Bayesian update to guarantee differential privacy and ensure consensus agreement and asymptotic learning for the agents.

Privacy-preserving methods in healthcare. Privacy and data security are critical in healthcare contexts (ARPA-H, 2024). An example is multicenter clinical trials, which are often difficult to coordinate and conduct; see, for example, the AIDS Clinical Trials Group (Campbell et al., 2012; Hammer et al., 1996). In the literature, a variety of methods and protocols have been proposed for privacy-preserving data sharing between collaborating healthcare centers and data providers; see, for example, Wan et al. (2022), Rehm et al. (2021), Froelicher et al. (2021), and Munjal and Bhatia (2023). Most of these methods are based on holomorphic encryption schemes and rely on centralized processing authority; see, e.g., Froelicher et al. (2021), Munjal and Bhatia (2023), and Geva et al. (2023)). In contrast, our method allows for fully distributed computation and does not require a centralized party to coordinate communication or computations between the centers. Furthermore, our method is based on statistical disclosure control following the differential privacy framework, which offers significant advantages in runtime and flexibility to accommodate a variety of statistical

models, including survival analysis (Geva et al., 2023; Froelicher et al., 2021), logistic regression (Geva et al., 2023), and linear regression (Froelicher et al., 2021).

Differentially private hypothesis testing. There have recently been several works on differentially private hypothesis testing (Canonne et al., 2019; Cai et al., 2017; Nissim et al., 2007). A work related to ours is the work of (Awan and Slavković, 2018), which devises the most powerful differentially private hypothesis test for binary variables, which introduces a new type of DP noise: the TULAP noise. Our work differs from the work of Awan and Slavković (2018) in two main axes: first, Awan and Slavković (2018) can be viewed as a centralized test, while ours is decentralized and can only handle binary signals, where our work can support arbitrary distributions. The work of Kazan et al. (2023) utilizes and adapts the mechanism Awan and Slavković (2018) to general distributions; however, the new mechanism is not distributed and does not have optimal statistical power. Moreover, the work of Canonne et al. (2019) shows that one of the noise distributions that optimize the sample complexity in non-distributed differentially private hypothesis testing is the Laplace noise. Similarly, in our work, we leverage the Laplace noise mechanism, which also has the added advantage of minimizing the convergence time of our distributed estimators.

Privacy in social network contexts. With our DP formulation for distributed learning environments, we can focus on privacy leaks through the information flow on the network rather than the underlying network structure. More broadly, the fundamental relationship between information flow and privacy is well noted, cf. contextual integrity (Nissenbaum, 2009; Benthall and Cummings, 2022) or how social network contexts affect link formation and information sharing behaviors (Acemoglu et al., 2022; Acemoglu et al., 2017). Notwithstanding, the privacy implications of information diffusion and algorithmic intervention on social networks are largely unexplored, except in a few studies (Rahimian et al., 2023b; Liell-Cock et al., 2020; Rahimian et al., 2023a). The existing work looks at this issue in specific, highly stylized scenarios. One study shows that it is extremely hard to hide the existence of a giant connected component of infected nodes under an independent cascade model, and a simple inference attack can reveal the status of a good fraction of nodes (Rezaei et al., 2021). In another study, the authors propose a decaying pseudo-noise added locally to mask the influence structure against an external observer with access to the entire information flow under the Friedkin-Johnsen influence propagation model (Liell-Cock et al., 2020).

Literature on distributed learning and estimation. Our work builds upon the distributed learning algorithms presented in Rahimian and Jadbabaie (2016b), by incorporating DP. Even though the belief updates posited in Rahimian and Jadbabaie (2016b) are similar to the ones presented in this paper, the incorporation of noise makes the algorithms' behavior significantly different, which requires novel tools to analyze that the current paper provides. Moreover, the work of Rahimian and Jadbabaie (2016b) provides asymptotic results for online learning and estimation, whereas, in our

paper, we provide nonasymptotic results which flesh out the trade-offs between the privacy budget, the number of agents, the network, the errors, and the statistical models of the agents.

Moreover, an essential distinction between estimating a continuous quantity in Papachristou and Rahimian (2024), Li et al. (2018), Kang et al. (2023), Rizk et al. (2023), and Youn et al. (2023) and choosing from a discrete set in this work is the fact that DP randomization in our work induces a non-trivial failure probability that is bounded away from zero with increasing number of iterations. Hence, in our work, to satisfactorily control the quality of the group decision outcome, we need the agents to repeat their deliberations in multiple rounds and then aggregate the outcome of several rounds towards a collective belief. Different aggregation schemes are possible (e.g., based on arithmetic or geometric averaging of the beliefs and using different threshold rules to identify the selected alternatives), which lead to different false-positive and missed detection rates and come with their own privacy guarantee and communication complexity requirements.

Appendix B: Non-Private Belief Propagation Algorithms

B.1. Non-private Distributed MLE

In the non-private regime, agents form beliefs $\boldsymbol{\mu}_{i,t}(\hat{\theta})$ over the states $\hat{\theta} \in \Theta$ at every iteration t and exchange their beliefs with their neighbors until they can detect the MLE. Rahimian and Jadbabaie (2016b) give Algorithm 4 for belief exchange and prove that beliefs converge to a uniform distribution over the MLE set (see Theorem 1 in their paper).

Algorithm 4 Non-Private Distributed MLE

The agents begin by forming: $\boldsymbol{\gamma}_i(\hat{\theta}) = \prod_{j=1}^{n_i} \ell_i(\boldsymbol{s}_j|\hat{\theta})$, and initializing their beliefs to $\boldsymbol{\mu}_{i,0}(\hat{\theta}) = \boldsymbol{\gamma}_i(\hat{\theta}) / \sum_{\tilde{\theta}\in\Theta} \boldsymbol{\gamma}_i(\tilde{\theta})$. In any future time period, the agents update their belief after communication with their neighboring agents, and according to the following update rule:

$$\boldsymbol{\mu}_{i,t}(\hat{\theta}) = \frac{\boldsymbol{\mu}_{i,t-1}^{1+a_{ii}}(\hat{\theta}) \prod_{j\in\mathcal{N}_i} \boldsymbol{\mu}_{j,t-1}^{a_{ij}}(\hat{\theta})}{\sum_{\tilde{\theta}\in\Theta} \boldsymbol{\mu}_{i,t-1}^{1+a_{ii}}(\tilde{\theta}) \prod_{j\in\mathcal{N}_i} \boldsymbol{\mu}_{j,t-1}^{a_{ij}}(\tilde{\theta})}, \text{for all } \hat{\theta}\in\Theta \text{ and } t\in[T]. \tag{18}$$

The convergence of this algorithm relies on studying the behavior of the log-belief ratio between two states $\hat{\theta}$ and $\check{\theta}$, whose dynamics are governed by the powers of the primitive matrix $A+I$, see (Rahimian and Jadbabaie, 2016b, Theorem 1), with modulus-ordered eigenvalues $0 < |\lambda_n(A+I)| \leq \cdots \leq |\lambda_2(A+I)| < \lambda_1(A+I) = 2$. Specifically, Rahimian and Jadbabaie (2016b, Theorem 1) state that there is a dominant term due to the maximum eigenvalue $\lambda_1(A+I) = 2$ that dominates the log-belief ratios for large t, and $n-1$ terms that decay exponentially with rate $(|\lambda_2(A+I)|/2)^t$. The dominant term depends on the difference of the log-likelihoods between the states $\hat{\theta}$ and $\check{\theta}$, namely $\Lambda(\hat{\theta},\check{\theta}) = \Lambda(\hat{\theta}) - \Lambda(\check{\theta})$. Subsequently, the log-belief ratio approaches $-\infty$ as $t\to\infty$ whenever $\check{\theta}\in\Theta^\star$, and we can recover the MLEs.

B.2. Non-private Online Learning from Intermittent Streams

Rahimian and Jadbabaie (2016b) give Algorithm 5 for the agents to determine the true state θ° from their streams of observations in the online learning regime.

Algorithm 5 Non-Private Distributed Online Learning

Every time $t \in \mathbb{N}_0$, each agent forms the likelihood product of the signals that it has received at that iteration: $\boldsymbol{\gamma}_{i,t}(\hat{\theta}) = \prod_{j=1}^{n_{i,t}} \ell_i(\boldsymbol{s}_{i,t}^j|\hat{\theta})$, if $n_{i,t} \geq 1$, and $\gamma_{i,t}(\hat{\theta}) = 1$ if $n_{i,t} = 0$. The agent then updates its belief according to:

$$\boldsymbol{\mu}_{i,t}(\hat{\theta}) = \frac{\boldsymbol{\gamma}_{i,t}(\hat{\theta})\boldsymbol{\mu}_{i,t-1}^{a_{ii}}(\hat{\theta}) \prod_{j\in\mathcal{N}_i} \boldsymbol{\mu}_{j,t-1}^{a_{ij}}(\hat{\theta})}{\sum_{\tilde{\theta}\in\Theta} \boldsymbol{\gamma}_{i,t}(\tilde{\theta})\boldsymbol{\mu}_{i,t-1}^{a_{ii}}(\tilde{\theta}) \prod_{j\in\mathcal{N}_i} \boldsymbol{\mu}_{j,t-1}^{a_{ij}}(\tilde{\theta})}, \text{for all } \hat{\theta} \in \Theta \text{ and } t \in [T], \tag{19}$$

initialized by: $\boldsymbol{\mu}_{i,0}(\hat{\theta}) = \boldsymbol{\gamma}_{i,0}(\hat{\theta}) / \sum_{\tilde{\theta}\in\Theta} \boldsymbol{\gamma}_{i,0}(\tilde{\theta})$.

Rahimian and Jadbabaie (2016b) prove that this algorithm converges to learning the true state asymptotically. Their argument relies on the fact that the time average of the log-belief ratio between any state $\hat{\theta} \in \Theta$ and the true state θ°, $(1/t)\log(\boldsymbol{\mu}_{i,t}(\hat{\theta})/\boldsymbol{\mu}_{i,t}(\theta^\circ))$, converges to the weighted sum of the KL divergences of all agents, which is less than zero if the models are statistically identifiable.

Appendix C: Proofs and Convergence Analysis for Distributed, Private MLE

C.1. Log-Belief Ratio Notations

Let $\boldsymbol{\mu}_{i,t}(\hat{\theta})$ be the beliefs for the nonprivate system, and let $\boldsymbol{\nu}_{i,k,t}(\hat{\theta})$ be the private estimates for agent $i \in [n]$, at the time step $t \geq 0$ of round $k \in [K]$. For the non-private algorithm (Algorithm 4) all pairs $\hat{\theta}, \check{\theta} \in \Theta$ we let $\boldsymbol{\phi}_{i,t}(\hat{\theta}, \check{\theta}) = \log\left(\frac{\boldsymbol{\mu}_{i,t}(\hat{\theta})}{\boldsymbol{\mu}_{i,t}(\check{\theta})}\right)$ and $\boldsymbol{\lambda}_i(\hat{\theta}, \check{\theta}) = \log\left(\frac{\gamma_i(\hat{\theta})}{\gamma_i(\check{\theta})}\right)$ be the log belief ratios and the log-likelihood ratios of the agent $i \in [n]$ in the time step $t \geq 0$. For the private algorithm (Algorithm 1) we let $\boldsymbol{\psi}_{i,k,t}(\hat{\theta}, \check{\theta}) = \log\left(\frac{\boldsymbol{\nu}_{i,k,t}(\hat{\theta})}{\boldsymbol{\nu}_{i,k,t}(\check{\theta})}\right)$, $\boldsymbol{\zeta}_{i,k}(\hat{\theta}, \check{\theta}) = \log\left(\frac{\boldsymbol{\sigma}_{i,k}(\hat{\theta})}{\boldsymbol{\sigma}_{i,k}(\check{\theta})}\right)$, and $\boldsymbol{\kappa}_{i,k}(\hat{\theta}, \check{\theta}) = \boldsymbol{d}_{i,k}(\hat{\theta}) - \boldsymbol{d}_{i,k}(\check{\theta})$ be the private log-belief ratio, log-likelihood ratio and log of the noise ratio between states $\hat{\theta}$ and $\check{\theta}$ for agent $i \in [n]$ at time step $t \geq 0$ of round $k \in [K]$, and denote their vectorized versions by $\overline{\boldsymbol{\phi}}_t(\hat{\theta}, \check{\theta})$, $\overline{\boldsymbol{\psi}}_{k,t}(\hat{\theta}, \check{\theta})$, $\overline{\boldsymbol{\lambda}}(\hat{\theta}, \check{\theta})$, and $\overline{\boldsymbol{\zeta}}_k(\hat{\theta}, \check{\theta})$, respectively, where the vectorization is over the agents $i \in [n]$. We define $\Lambda(\hat{\theta}, \check{\theta}) = \mathbb{1}^T \overline{\boldsymbol{\lambda}}(\hat{\theta}, \check{\theta})$, $Z_k(\hat{\theta}, \check{\theta}) = \mathbb{1}^T \overline{\boldsymbol{\zeta}}_k(\hat{\theta}, \check{\theta}) = \mathbb{1}^T \left(\overline{\boldsymbol{\lambda}}(\hat{\theta}, \check{\theta}) + \overline{\boldsymbol{\kappa}}_k(\hat{\theta}, \check{\theta})\right)$.

We say that a stochastic process $\{X_t\}_{t\in\mathbb{N}}$ converges in L_2 to X, and write $X_t \xrightarrow{L_2} X$, if and only if $\lim_{t\to\infty} \mathbb{E}\left[\|X_t - X\|_2\right] = 0$. Convergence in L_2 implies convergence in probability (i.e., $X_t \xrightarrow{p} X$) due to Markov's inequality.

To prove the results, we need the following auxiliary lemma:

Lemma 1. *Let $X_1, \ldots, X_n$ be i.i.d. draws from a distribution $\mathcal{D}$. Then, for every $i \in [n]$, we have $\mathbb{P}[X_i \geq \max_{j\neq i} X_j] = 1/n$.*

Let $E_i = \{X_i \geq \max_{j\neq i} X_j\}$. Note that because $X_1, \ldots, X_n$ are i.i.d. $\mathbb{P}[E_1] = \mathbb{P}[E_2] = \cdots = \mathbb{P}[E_n]$. Moreover, note that since the maximum is unique, we have that $E_1, \ldots, E_n$ are a partition of the sample space. Therefore, we get $\sum_{i\in[n]} \mathbb{P}[E_i] = 1$, and subsequently $\mathbb{P}[E_i] = 1/n$. Q.E.D.

C.2. Proof of Theorem 1: Asymptotic Behavior of the AM/GM Algorithm

Similarly to Rahimian and Jadbabaie, 2016b, Theorem 1, we have that for any $k \in [K]$

$$\begin{aligned}\left\|\overline{\boldsymbol{\psi}}_{k,t}(\hat{\theta},\check{\theta}) - \frac{2^t}{n} Z_k(\hat{\theta},\check{\theta})\mathbb{1}\right\|_2 &= \left\|\sum_{i=2}^{n}\left(\frac{\lambda_i(A+I)}{2}\right)^t \bar{\boldsymbol{l}}_i \overline{\boldsymbol{r}}_i^T \overline{\boldsymbol{\zeta}}_k(\hat{\theta},\check{\theta})\right\|_2 \\ &\le \sum_{i=2}^{n}\left|\frac{\lambda_i(A+I)}{2}\right|^t \|\bar{\boldsymbol{l}}_i\|_2 \left|\overline{\boldsymbol{r}}_i^T \overline{\boldsymbol{\zeta}}_k(\hat{\theta},\check{\theta})\right| \\ &\le \sum_{i=2}^{n}\left|\frac{\lambda_i(A+I)}{2}\right|^t \left\|\overline{\boldsymbol{\zeta}}_k(\hat{\theta},\check{\theta})\right\|_2 \\ &\le \sum_{i=2}^{n}\left|\frac{\lambda_i(A+I)}{2}\right|^t \left(\left\|\overline{\boldsymbol{\lambda}}(\hat{\theta},\check{\theta})\right\|_2 + \left\|\overline{\boldsymbol{\kappa}}_k(\hat{\theta},\check{\theta})\right\|_2\right).\end{aligned}$$

Taking expectations and applying Jensen's inequality, we can bound the above as

$$\mathbb{E}\left[\left\|\overline{\boldsymbol{\psi}}_{k,t}(\hat{\theta},\check{\theta}) - \frac{2^t}{n} Z_k(\hat{\theta},\check{\theta})\mathbb{1}\right\|_2\right] \le 2(n-1)\left|\frac{\lambda_2(A+I)}{2}\right|^t \left[n\Gamma_{n,\Theta} + \sum_{i=1}^{n}\sqrt{\mathbb{V}\left[\boldsymbol{d}_{i,k}(\hat{\theta})\right]}\right], \tag{20}$$

where $\Gamma_{n,\Theta} = \max_{i\in[n],\hat{\theta}\in\Theta} |\log \boldsymbol{\gamma}_i(\hat{\theta})|$ and $|\lambda_2(A+I)|/2 < 1$ is the SLEM of $(A+I)/2$, henceforth also denoted by $a_n^\star$; A is doubly stochastic so $0 < |\lambda_n(A+I)| \le |\lambda_{n-1}(A+I)| \le \cdots \le |\lambda_2(A+I)| < \lambda_1(A+I) = 2$. Note that the right-hand side in Equation (20) goes to 0 as $t \to \infty$, which implies that (by Markov's inequality) $\overline{\boldsymbol{\psi}}_{k,t}(\hat{\theta},\check{\theta}) \xrightarrow{L_2} \frac{2^t}{n} Z_k(\hat{\theta},\check{\theta})\mathbb{1}$.

Let $\theta^\star \in \Theta^\star$ and $\bar{\theta} \in \bar{\Theta}$. Based on our definition of $\Lambda(\cdot,\cdot)$ we should have $\Lambda(\bar{\theta},\theta^\star) < 0$ and the corresponding non-private algorithm would have $\phi_{i,t}(\bar{\theta},\theta^\star) \to -\infty$ for every $\bar{\theta} \in \bar{\Theta}$ and $\theta^\star \in \Theta^\star$. However, when noise is introduced, there could be mistakes introduced in the log-belief ratios, even if $\Lambda(\bar{\theta},\theta^\star) < 0$ we can have $Z_k(\bar{\theta},\theta^\star) \ge 0$ which implies $\boldsymbol{\psi}_{i,k,t}(\bar{\theta},\theta^\star) \not\to -\infty$.

AM estimator. For all $\hat{\theta} \in \Theta$ we let $\boldsymbol{Y}_j(\hat{\theta}) = \sum_{i=1}^{n} \boldsymbol{d}_{i,k}(\hat{\theta})$. We focus on a single $\theta^\star \in \Theta^\star$. It is easy to show that because the noise is independent across the agents, and i.i.d. between the states for a given agent i, then $\boldsymbol{Y}_j(\bar{\theta})$ and $\boldsymbol{Y}_j(\theta^\star)$ are also i.i.d. for all $\bar{\theta} \in \bar{\Theta}$, and therefore we have that (see Lemma 1)

$$\lim_{T\to\infty} \mathbb{P}\left[\boldsymbol{\nu}_{i,k,T}(\theta^\star) > 0\right] \ge \mathbb{P}\left[\bigcap_{\bar{\theta}\in\bar{\Theta}} \left\{\boldsymbol{Y}_k(\bar{\theta}) \le \boldsymbol{Y}_k(\theta^\star)\right\}\right] = \frac{1}{|\bar{\Theta}|}.$$

Therefore, for the AM estimator we have the following:

$$\lim_{T\to\infty} \mathbb{P}\left[\boldsymbol{\nu}_{i,T}^{\mathrm{AM}}(\theta^\star) = 0\right] = \lim_{T\to\infty} \mathbb{P}\left[\bigcap_{k\in[K]} \left\{\boldsymbol{\nu}_{i,k,T}(\theta^\star) = 0\right\}\right] \le \left(1 - \frac{1}{|\bar{\Theta}|}\right)^K \le e^{-\frac{K}{|\bar{\Theta}|}}, \tag{21}$$

and the failure probability of the AM estimator can be calculated by applying the union bound as

$$\begin{aligned}\lim_{T\to\infty} \mathbb{P}\left[\Theta^\star \not\subseteq \hat{\Theta}_{i,T}^{\mathrm{AM}}\right] &= \lim_{T\to\infty} \mathbb{P}\left[\exists \theta^\star \in \Theta^\star : \boldsymbol{\nu}_{i,T}^{\mathrm{AM}}(\theta^\star) = 0\right] \\ &\le \lim_{T\to\infty} \sum_{\theta^\star\in\Theta^\star} \mathbb{P}\left[\boldsymbol{\nu}_{i,T}^{\mathrm{AM}}(\theta^\star) = 0\right] \\ &\le |\Theta^\star| e^{-\frac{K}{|\bar{\Theta}|}}.\end{aligned}$$

Setting $K = |\bar{\Theta}| \log(|\Theta^\star|/(1-\beta))$, we can ensure that the Type II error rate is at most $1-\beta$.

GM estimator. Similarly, for $\boldsymbol{\nu}_{i,T}^{\mathrm{GM}}$ we have the following:

$$\lim_{T\to\infty} \mathbb{P}\left[\boldsymbol{\nu}_{i,k,T}(\bar{\theta}) = 0\right] \geq \mathbb{P}\left[\bigcap_{\theta^\star\in\Theta^\star} \left\{\boldsymbol{Y}_k(\bar{\theta}) \leq \boldsymbol{Y}_k(\theta^\star)\right\}\right] = \frac{1}{|\Theta^\star|},$$

for all $\bar{\theta} \in \bar{\Theta}$ and subsequently

$$\lim_{T\to\infty} \mathbb{P}\left[\boldsymbol{\nu}_{i,T}^{\mathrm{GM}}(\bar{\theta}) > 0\right] = \lim_{T\to\infty} \mathbb{P}\left[\bigcap_{k\in[K]} \left\{\boldsymbol{\nu}_{i,k,T}(\bar{\theta}) > 0\right\}\right] \leq \left(1 - \frac{1}{|\Theta^\star|}\right)^K \leq e^{-\frac{K}{|\Theta^\star|}},$$

and

$$\begin{aligned}
\lim_{T\to\infty} \mathbb{P}\left[\hat{\Theta}_{i,T}^{\mathrm{GM}} \not\subseteq \Theta^\star\right] &= \lim_{T\to\infty} \mathbb{P}\left[\hat{\Theta}_{i,T}^{\mathrm{GM}} \cap \bar{\Theta} \neq \emptyset\right] \\
&= \lim_{T\to\infty} \mathbb{P}\left[\exists \bar{\theta} \in \bar{\Theta} : \boldsymbol{\nu}_{i,T}^{\mathrm{GM}}(\bar{\theta}) > 0\right] \\
&\leq \lim_{T\to\infty} \sum_{\bar{\theta}\in\bar{\Theta}} \mathbb{P}\left[\boldsymbol{\nu}_{i,T}^{\mathrm{GM}}(\bar{\theta}) > 0\right] \\
&\leq |\bar{\Theta}| e^{-\frac{K}{|\Theta^\star|}}.
\end{aligned}$$

Therefore, selecting $K = |\Theta^\star| \log(|\bar{\Theta}|/\alpha)$ produces a Type I error rate of at most α.

Differentially Private Outputs. Due to the postprocessing property of DP (see Proposition 2.1 of Dwork and Roth (2014)), the resulting estimates are ε-DP with respect to the private signals.

C.3. Proof of Theorem 2: Finite-Time Guarantees for AM/GM Algorithm

For simplicity of exposition, we have set $\varrho^{\mathrm{AM}} = \varrho^{\mathrm{GM}} = \varrho$ which corresponds to $\tau^{\mathrm{AM}} = \tau^{\mathrm{GM}} = \tau$.

GM estimator. We let $\boldsymbol{Y}_k(\hat{\theta}) = \sum_{i\in[n]} \boldsymbol{d}_{i,k}(\hat{\theta})$. We define the event $F = \left\{\exists k \in [K], \bar{\theta} \in \bar{\Theta}, \theta^\star \in \Theta^\star : Z^{\mathrm{GM}}(\bar{\theta}, \theta^\star) > \Lambda(\bar{\theta}, \theta^\star)\right\}$. We have that

$$\begin{aligned}
\mathbb{P}[F] &\overset{\text{(i)}}{\leq} \sum_{\bar{\theta}\in\bar{\Theta}} \mathbb{P}\left[\exists k \in [K], \theta^\star \in \Theta^\star : Z_k(\bar{\theta}, \theta^\star) > \Lambda(\bar{\theta}, \theta^\star)\right] \\
&\overset{\text{(ii)}}{=} \sum_{\bar{\theta}\in\bar{\Theta}} \left(1 - \mathbb{P}\left[\forall \theta^\star \in \Theta^\star : Z_1(\bar{\theta}, \theta^\star) \leq \Lambda(\bar{\theta}, \theta^\star)\right]\right)^K \\
&\overset{\text{(iii)}}{\leq} |\bar{\Theta}| \left(1 - \mathbb{P}\left[\bigcap_{\bar{\theta}\in\bar{\Theta}} \left\{\boldsymbol{Y}_1(\bar{\theta}) \leq \boldsymbol{Y}_1(\theta^\star)\right\}\right]\right)^K \\
&\overset{\text{(iv)}}{\leq} |\bar{\Theta}| \left(1 - \frac{1}{|\Theta^\star|}\right)^K \\
&\overset{\text{(v)}}{\leq} |\bar{\Theta}| e^{-K/|\Theta^\star|},
\end{aligned}$$

where the above follows (i) from the application of the union bound, (ii) using the fact that the K rounds are independent, (iii) the fact that $\{Z_1(\bar{\theta}, \theta^\star) \leq \Lambda(\bar{\theta}, \theta^\star)\} \supseteq \bigcap_{\bar{\theta}\in\bar{\Theta}} \left\{\boldsymbol{Y}_k(\bar{\theta}) \leq \boldsymbol{Y}_k(\theta^\star)\right\}$ and that $\{\boldsymbol{Y}_k(\bar{\theta})\}_{\bar{\theta}\in\bar{\Theta}}$ and $\boldsymbol{Y}_k(\theta^\star)$ are i.i.d., (iv) Lemma 1, and (v) $1 + x \leq e^x$ for all $x \in \mathbb{R}$.

For each round k of the algorithm and each pair of states $\hat{\theta}, \check{\theta} \in \Theta$, the following holds for the log-belief ratio of the geometrically averaged estimates:

$$\boldsymbol{\psi}_{i,t}^{\mathrm{GM}}(\hat{\theta}, \check{\theta}) = \frac{1}{K} \sum_{k \in [K]} \boldsymbol{\psi}_{i,k,t}(\hat{\theta}, \check{\theta}).$$

We let $Z^{\mathrm{GM}}(\hat{\theta}, \check{\theta}) = \frac{1}{K} \sum_{k=1}^{K} Z_k(\hat{\theta}, \check{\theta})$. By the triangle inequality and Theorem 1

$$\mathbb{E}\left[\left\|\overline{\boldsymbol{\psi}}_t(\hat{\theta}, \check{\theta}) - \frac{2^t}{n} Z^{\mathrm{GM}}(\hat{\theta}, \check{\theta}) \mathbb{1}\right\|_2\right] \leq \frac{1}{\sqrt{K}} \sum_{k=1}^{K} \mathbb{E}\left[\left\|\overline{\boldsymbol{\psi}}_{k,t}(\hat{\theta}, \check{\theta}) - \frac{2^t}{n} Z_k(\hat{\theta}, \check{\theta}) \mathbb{1}\right\|_2\right]. \tag{22}$$

For every round $k \in [K]$,

$$\begin{aligned}
\left\|\overline{\boldsymbol{\psi}}_{k,t}(\hat{\theta}, \check{\theta}) - \frac{2^t}{n} Z_k(\hat{\theta}, \check{\theta}) \mathbb{1}\right\|_2 &= \left\|\sum_{i=2}^{n} \left(\frac{\lambda_i(I+A)}{2}\right)^t \bar{\boldsymbol{l}}_i \overline{\boldsymbol{r}}_i^T \overline{\boldsymbol{\zeta}}_k(\hat{\theta}, \check{\theta})\right\|_2 \\
&\leq \sum_{i=2}^{n} \left|\frac{\lambda_i(I+A)}{2}\right|^t \|\bar{\boldsymbol{l}}_i\|_2 \left|\overline{\boldsymbol{r}}_i^T \overline{\boldsymbol{\zeta}}_k(\hat{\theta}, \check{\theta})\right| \\
&\leq \sum_{i=2}^{n} \left|\frac{\lambda_i(I+A)}{2}\right|^t \left\|\overline{\boldsymbol{\zeta}}_k(\hat{\theta}, \check{\theta})\right\|_2 \\
&\leq \sum_{i=2}^{n} \left|\frac{\lambda_i(I+A)}{2}\right|^t \left(\left\|\overline{\boldsymbol{\lambda}}(\hat{\theta}, \check{\theta})\right\|_2 + \left\|\overline{\boldsymbol{\kappa}}_k(\hat{\theta}, \check{\theta})\right\|_2\right).
\end{aligned}$$

Taking expectations and applying Jensen's inequality, we can bound the above as

$$\mathbb{E}\left[\left\|\overline{\boldsymbol{\psi}}_{k,t}(\hat{\theta}, \check{\theta}) - \frac{2^t}{n} Z_k(\hat{\theta}, \check{\theta}) \mathbb{1}\right\|_2\right] \leq 2(n-1) \left|\frac{\lambda_2(I+A)}{2}\right|^t \left[n\Gamma_{n,\Theta} + \sum_{i=1}^{n} \sqrt{\mathbb{V}\left[\boldsymbol{d}_{i,k}(\hat{\theta})\right]}\right]. \tag{23}$$

Using $a_n^\star = \left|\frac{\lambda_2(I+A)}{2}\right|$, and combining Equation (22) with Equation (23), we get:

$$\mathbb{E}\left[\left\|\overline{\boldsymbol{\psi}}_t(\hat{\theta}, \check{\theta}) - \frac{2^t}{n} Z^{\mathrm{GM}}(\hat{\theta}, \check{\theta}) \mathbb{1}\right\|_2\right] \leq \frac{2(n-1)(a_n^\star)^t [n\Gamma_{n,\Theta} + V_{n,\Theta}]}{\sqrt{K}}. \tag{24}$$

By Markov's inequality and Equation (22), we get that for every $z > 0$

$$\mathbb{P}\left[\left\|\overline{\boldsymbol{\psi}}_t^{\mathrm{GM}}(\hat{\theta}, \check{\theta}) - \frac{2^t}{n} Z^{\mathrm{GM}}(\hat{\theta}, \check{\theta}) \mathbb{1}\right\|_2 > z\right] \leq \frac{2(n-1)(a_n^\star)^t [n\Gamma_{n,\Theta} + V_{n,\Theta}]}{z\sqrt{K}}.$$

We let the RHS be equal to $\alpha/(|\Theta^\star||\bar{\Theta}|)$, which corresponds to letting $z = |\bar{\Theta}||\Theta^\star| \frac{2(n-1)(a_n^\star)^t [n\Gamma_{n,\Theta} + V_{n,\Theta}]}{\alpha\sqrt{K}}$. By applying a union bound over $\bar{\Theta}$ and $\Theta^\star$ we have that for every $\bar{\theta} \in \bar{\Theta}$ and $\theta^\star \in \Theta^\star$, with probability at least $1 - \alpha$,

$$\boldsymbol{\psi}_{i,t}^{\mathrm{GM}}(\bar{\theta}, \theta^\star) \leq \frac{2^t}{n} Z^{\mathrm{GM}}(\bar{\theta}, \theta^\star) + |\bar{\Theta}||\Theta^\star| \frac{2(n-1)(a_n^\star)^t [n\Gamma_{n,\Theta} + V_{n,\Theta}]}{\alpha\sqrt{K}}. \tag{25}$$

Conditioned on F^c, the above becomes

$$\begin{aligned}
\boldsymbol{\psi}_{i,t}^{\mathrm{GM}}(\bar{\theta}, \theta^\star) &\leq \frac{2^t}{n} \Lambda(\bar{\theta}, \theta^\star) + |\bar{\Theta}||\Theta^\star| \frac{2(n-1)(a_n^\star)^t [n\Gamma_{n,\Theta} + V_{n,\Theta}]}{\alpha\sqrt{K}} \\
&\leq -\frac{2^t}{n} l_{n,\Theta} + |\Theta|^2 \frac{(n-1)(a_n^\star)^t [n\Gamma_{n,\Theta} + V_{n,\Theta}]}{2\alpha\sqrt{K}}.
\end{aligned} \tag{26}$$

Note that we have used $|\bar{\Theta}||\Theta^\star| = (|\Theta| - |\Theta^\star|)|\Theta^\star| \leq |\Theta|^2/4$. To make Equation (26) at most $-\varrho$ for some $\varrho > 0$, we need to set

$$t \geq \max\left\{\frac{\log\left(\frac{2\varrho n}{l_{n,\Theta}}\right)}{\log 2}, \frac{\log\left(\frac{|\Theta|^2(n-1)(n\Gamma_{n,\Theta}+V_{n,\Theta})}{2\alpha\varrho\sqrt{K}}\right)}{\log(1/a_n^\star)}\right\} = T.$$

The log-belief ratio threshold implies that for all $\theta^\star \in \Theta^\star, \bar{\theta} \in \bar{\Theta}$ we have that $\boldsymbol{\nu}_{i,k,T}(\theta^\star) \geq e^\varrho \boldsymbol{\nu}_{i,k,T}(\bar{\theta})$. To determine the belief threshold τ, note that

$$\begin{aligned} 1 &= \sum_{\hat{\theta}\in\hat{\Theta}_{i,T}^{\mathrm{GM}}} \boldsymbol{\nu}_{i,k,T}(\hat{\theta}) + \sum_{\hat{\theta}\notin\hat{\Theta}_{i,T}^{\mathrm{GM}}} \boldsymbol{\nu}_{i,k,T}(\hat{\theta}) \geq (1+e^\varrho) \max_{\hat{\theta}\notin\Theta_{i,T}^{\mathrm{GM}}} \boldsymbol{\nu}_{i,k,T}(\hat{\theta}) \\ &\implies \max_{\hat{\theta}\notin\Theta_{i,T}^{\mathrm{GM}}} \boldsymbol{\nu}_{i,k,T}(\hat{\theta}) \leq \frac{1}{1+e^\varrho}. \end{aligned}$$

Moreover, we can prove that $\min_{\hat{\theta}\in\hat{\Theta}_{i,T}^{\mathrm{GM}}} \boldsymbol{\nu}_{i,k,T}(\hat{\theta}) \geq \frac{1}{1+e^{-\varrho}} \geq \frac{1}{1+e^\varrho}$ since $\varrho > 0$ which shows that any value in $[1/(1+e^\varrho), 1/(1+e^{-\varrho})]$ is a valid threshold. This yields that $\mathbb{P}[\hat{\Theta}_{i,T}^{\mathrm{GM}} \subseteq \Theta^\star | F^c] = \mathbb{P}[\hat{\Theta}_{i,T}^{\mathrm{GM}} \cap \bar{\Theta} = \emptyset | F^c] = \mathbb{P}[\exists \bar{\theta} \in \bar{\Theta} : \boldsymbol{\nu}_{i,T}^{\mathrm{GM}}(\bar{\theta}) > \tau | F^c] \geq 1-\alpha$. Subsequently, by letting $K = |\Theta^\star| \log(|\bar{\Theta}|/\alpha)$, we get that

$$\mathbb{P}[\hat{\Theta}_{i,T}^{\mathrm{GM}} \subseteq \Theta^\star] \geq \mathbb{P}[\hat{\Theta}_{i,T}^{\mathrm{GM}} \subseteq \Theta^\star | F^c]\mathbb{P}[F^c] \geq (1-\alpha)^2 \geq 1-2\alpha.$$

AM estimator. We let $E = \left\{\exists k \in [K], \bar{\theta} \in \bar{\Theta}, \theta^\star \in \Theta^\star : Z_k(\theta^\star, \bar{\theta}) < \Lambda(\theta^\star, \bar{\theta})\right\}$. Using similar arguments to Theorem 2, we can deduce that $\mathbb{P}[E] \leq |\Theta^\star| e^{-K/|\bar{\Theta}|}$. By setting $K = |\bar{\Theta}| \log(|\Theta^\star|/(1-\beta))$, we make $\mathbb{P}[E] \leq 1-\beta$.

Conditioned on E^c and by applying Markov's inequality, similarly to Equation (26) for the GM estimator, we have that for all $\theta^\star \in \Theta^\star, \bar{\theta} \in \bar{\Theta}$ and run $k \in [K]$

$$\begin{aligned} \boldsymbol{\psi}_{i,k,t}(\theta^\star, \bar{\theta}) &\geq \frac{2^t}{n} Z_k(\theta^\star, \bar{\theta}) - |\Theta^\star||\bar{\Theta}| \frac{2(n-1)(a_n^\star)^t(n\Gamma_{n,\Theta}+V_{n,\Theta})}{1-\beta'} \\ &\geq \frac{2^t}{n}\Lambda(\theta^\star, \bar{\theta}) - |\Theta|^2 \frac{(n-1)(a_n^\star)^t(n\Gamma_{n,\Theta}+V_{n,\Theta})}{2(1-\beta')} \\ &\geq \frac{2^t}{n} l_{n,\Theta} - |\Theta|^2 \frac{2(n-1)(a_n^\star)^t(n\Gamma_{n,\Theta}+V_{n,\Theta})}{2(1-\beta')}, \end{aligned} \tag{27}$$

with probability β' for some $\beta' \in (0,1)$. To make the above at least ϱ for some $\varrho > 0$ it suffices to pick

$$t \geq \max\left\{\frac{\log\left(\frac{2\varrho n}{l_{n,\Theta}}\right)}{\log 2}, \frac{\log\left(\frac{|\Theta|^2(n-1)(n\Gamma_{n,\Theta}+V_{n,\Theta})}{2(1-\beta')\varrho}\right)}{\log(1/a_n^\star)}\right\} = T.$$

The log-belief ratio threshold implies that for all $\theta^\star \in \Theta^\star, \bar{\theta} \in \bar{\Theta}$ we have that $\boldsymbol{\nu}_{i,k,T}(\theta^\star) \geq e^\varrho \boldsymbol{\nu}_{i,k,T}(\bar{\theta})$. Similarly, any value in $[1/(1+e^\varrho), 1/(1+e^{-\varrho})]$ is a valid threshold, we can show that

$$
\begin{aligned}
\mathbb{P}[\Theta^\star \subseteq \hat{\Theta}^{\mathrm{AM}}_{i,T} | E^c] &= \mathbb{P}\left[\forall \theta^\star \in \Theta^\star : \boldsymbol{\nu}^{\mathrm{AM}}_{i,T}(\theta^\star) > \tau \middle| E^c\right] \\
&\geq \mathbb{P}\left[\bigcup_{k\in[K]} \{\forall \theta^\star \in \Theta^\star : \boldsymbol{\nu}_{i,k,T}(\theta^\star) > \tau\} \middle| E^c\right] \\
&= 1 - \mathbb{P}\left[\bigcap_{k\in[K]} \{\exists \theta^\star \in \Theta^\star : \boldsymbol{\nu}_{i,k,T}(\theta^\star) < \tau\} \middle| E^c\right] \\
&= 1 - \left(\mathbb{P}\left[\exists \theta^\star \in \Theta^\star : \boldsymbol{\nu}_{i,k,T}(\theta^\star) < \tau \middle| E^c\right]\right)^K \\
&= 1 - \left(1 - \mathbb{P}\left[\forall \theta^\star \in \Theta^\star : \boldsymbol{\nu}_{i,k,T}(\theta^\star) > \tau \middle| E^c\right]\right)^K \\
&= 1 - (\beta')^K \\
&\geq 1 - e^{-(1-\beta')K}.
\end{aligned}
$$

We set $1 - \beta' = \log(1/(1-\beta))/K$ and have that

$$
\mathbb{P}[\Theta^\star \subseteq \hat{\Theta}^{\mathrm{AM}}_{i,T}] \geq \mathbb{P}[E^c]\mathbb{P}[\Theta^\star \subseteq \hat{\Theta}^{\mathrm{AM}}_{i,T} | E^c] \geq (\beta)^2 \geq 1 - 2(1-\beta).
$$

These finally yield

$$
T = \max\left\{\frac{\log\left(\frac{2\varrho n}{l_{n,\Theta}}\right)}{\log 2}, \frac{\log\left(\frac{|\Theta|^2(n-1)K(n\Gamma_{n,\Theta}+V_{n,\Theta})}{2\log(1/(1-\beta))\varrho}\right)}{\log(1/a^\star_n)}\right\}.
$$

Privacy (for both estimators). Minimizing the convergence time T corresponds to minimizing $V_{n,\Theta}$. Since $V_{n,\Theta}$ is separable over the agents, it suffices to solve the problem of minimizing the variance of each noise variable independently.

If an adversary can eavesdrop only once and has access to any round $k \in [K]$, and we have $|\Theta|$ states, by the composition theorem, the budget ε should be divided by $K|\Theta|$. Thus, the problem of finding the optimal distribution $\mathcal{D}_i(\varepsilon)$ on $\mathbb{R}$, corresponds to the following optimization problem studied in Koufogiannis et al., 2015; Papachristou and Rahimian, 2024, for all $k \in [K]$:

$$
\begin{aligned}
\min_{\mathcal{D}_i(\varepsilon)\in\mathrm{Simplex}(\mathbb{R})} \quad & \mathbb{V}_{\boldsymbol{d}_{i,k}\sim\mathcal{D}_i(\varepsilon)}[\boldsymbol{d}_{i,k}] \\
\text{s.t.} \quad & \mathcal{D}_i(\varepsilon) \text{ is } \frac{\varepsilon}{K|\Theta|}\text{-DP and } \mathbb{E}_{\boldsymbol{d}_{i,k}\sim\mathcal{D}_i(\varepsilon)}[\boldsymbol{d}_{i,k}] = 0.
\end{aligned} \tag{28}
$$

The optimal solution to this problem is given by Koufogiannis et al. (2015) and Papachristou and Rahimian (2024). For privatizing beliefs in our problem, this corresponds to selecting $\mathcal{D}^\star_i(\varepsilon) = \mathrm{Lap}\left(\frac{\Delta_{n,\Theta}K|\Theta|}{\varepsilon}\right)$ where $\Delta_{n,\Theta}$ is the sensitivity of the log-likelihood of the i-th agent.

Differentially Private Outputs. Due to the postprocessing property of DP (see Proposition 2.1 of Dwork and Roth (2014)), the resulting estimates are ε-DP with respect to the private signals.

Q.E.D.

C.4. Proof of Theorem 3: Asymptotic Behavior of the Two-Threshold Algorithm

For simplicity, we define $p_2 = 1/|\bar{\Theta}|$ and $p_1 = 1 - 1/|\Theta^\star|$, and assume that $\varrho^{\text{thres},1} = \varrho^{\text{thres},2} = \varrho$ which implies that $\hat{\tau}^{\text{thres},1} = \hat{\tau}^{\text{thres},2} = \tau$ and $\boldsymbol{N}_{i,t}^{\text{thres},1}(\hat{\theta}) = \boldsymbol{N}_{i,t}^{\text{thres},2}(\hat{\theta}) = \boldsymbol{N}_{i,T}(\hat{\theta})$.

Type I estimator ($\hat{\Theta}_{i,T}^{\text{thres},1}$). To determine the asymptotic Type I error probability α of $\hat{\Theta}_{i,T}^{\text{thres},1}$, we assume that the threshold takes the form $\tau^{\text{thres},1} = (1+\pi_1)p_1$. Then

$$\begin{aligned}
\lim_{T\to\infty} \mathbb{P}\left[\hat{\Theta}_{i,T}^{\text{thres},1} \not\subseteq \Theta^\star\right] &\overset{\text{(i)}}{=} \lim_{T\to\infty} \mathbb{P}\left[\exists \bar{\theta} \in \bar{\Theta} : \bar{\theta} \in \hat{\Theta}_{i,T}^{\text{thres},1}\right] \\
&\overset{\text{(ii)}}{\leq} \lim_{T\to\infty} \sum_{\bar{\theta}\in\bar{\Theta}} \mathbb{P}\left[\bar{\theta} \in \hat{\Theta}_{i,T}^{\text{thres},1}\right] \\
&\overset{\text{(iii)}}{\leq} \lim_{T\to\infty} \sum_{\bar{\theta}\in\bar{\Theta}} \mathbb{P}\left[\boldsymbol{N}_{i,T}(\bar{\theta}) \geq \tau^{\text{thres},1}\right] \\
&\leq \lim_{T\to\infty} \sum_{\bar{\theta}\in\bar{\Theta}} \mathbb{P}\left[\boldsymbol{N}_{i,T}(\bar{\theta}) \geq (1+\pi_1)p_1\right] \\
&\overset{\text{(iv)}}{\leq} \lim_{T\to\infty} \sum_{\bar{\theta}\in\bar{\Theta}} \mathbb{P}\left[\boldsymbol{N}_{i,T}(\bar{\theta}) \geq (1+\pi_1)\mathbb{E}\left[N_{i,T}(\bar{\theta})\right]\right] \\
&\overset{\text{(v)}}{\leq} |\bar{\Theta}| e^{-2K\pi_1^2},
\end{aligned} \tag{29}$$

where the result is derived by applying (i) the definition of $\hat{\Theta}_{i,T}^{\text{thres},1}$, (ii) union bound, (iii) the definition of the threshold $\tau^{\text{thres},1} = (1+\pi_1)p_1$, (iv) the fact that $\mathbb{E}\left[\boldsymbol{N}_{i,T}(\bar{\theta})\right] \leq p_1$ for all $\bar{\theta} \in \bar{\Theta}$ as $T \to \infty$, and the Chernoff bound on $\boldsymbol{N}_{i,T}(\bar{\theta})$. Therefore, to make the above α, it suffices to choose $K = \frac{\log(|\bar{\Theta}|/\alpha)}{2\pi_1^2}$.

Type II estimator ($\hat{\Theta}_{i,T}^{\text{thres},2}$). To determine the asymptotic Type II error probability $1-\beta$ of $\hat{\Theta}_{i,T}^{\text{thres},2}$, we assume that the threshold takes the form $\tau^{\text{thres},2} = (1-\pi_2)p_2$. Then

$$\begin{aligned}
\lim_{T\to\infty} \mathbb{P}\left[\Theta^\star \not\subseteq \hat{\Theta}_{i,T}^{\text{thres},2}\right] &\overset{\text{(i)}}{=} \lim_{T\to\infty} \mathbb{P}\left[\exists \theta^\star \in \Theta^\star : \theta^\star \notin \Theta_{i,T}^{\text{thres},2}\right] \\
&\overset{\text{(ii)}}{\leq} \lim_{T\to\infty} \sum_{\theta^\star\in\Theta^\star} \mathbb{P}\left[\theta^\star \notin \Theta_{i,T}^{\text{thres},2}\right] \\
&\overset{\text{(iii)}}{\leq} \lim_{T\to\infty} \sum_{\theta^\star\in\Theta^\star} \mathbb{P}\left[\boldsymbol{N}_{i,T}(\theta^\star) \leq \tau^{\text{thres},2}\right] \\
&\overset{\text{(iii)}}{\leq} \lim_{T\to\infty} \sum_{\theta^\star\in\Theta^\star} \mathbb{P}\left[\boldsymbol{N}_{i,T}(\theta^\star) \leq (1-\pi_2)p_2\right] \\
&\overset{\text{(iv)}}{\leq} \lim_{T\to\infty} \sum_{\theta^\star\in\Theta^\star} \mathbb{P}\left[\boldsymbol{N}_{i,T}(\theta^\star) \leq (1-\pi_2)\mathbb{E}\left[\boldsymbol{N}_{i,T}(\theta^\star)\right]\right] \\
&\overset{\text{(v)}}{\leq} |\Theta^\star| e^{-2K\pi_2^2},
\end{aligned} \tag{30}$$

where the result is derived by applying (i) the definition of $\hat{\Theta}_{i,T}^{\text{thres},2}$, (ii) union bound, (iii) the definition of the threshold $\tau^{\text{thres},2} = (1-\pi_2)p_2$, (iv) the fact that $\mathbb{E}\left[\boldsymbol{N}_{i,T}(\theta^\star)\right] \geq p_2$ for all $\theta^\star \in \Theta^\star$ as $T \to \infty$, and (v) the Chernoff bound on $\boldsymbol{N}_{i,T}(\hat{\theta})$. Therefore, to make the above less than $1-\beta$, it suffices to choose $K = \frac{\log(|\bar{\Theta}|/(1-\beta))}{2\pi_2^2}$.

Differentially Private Outputs. Due to the immunity to post-processing of DP (Dwork and Roth, 2014, Proposition 2.1), the resulting estimates are ε-DP with respect to private signals.

Q.E.D.

C.5. Proof of Theorem 4: Finite-Time Guarantees for the Two-Threshold Algorithm

For simplicity, we assume that $\varrho^{\text{thres},1} = \varrho^{\text{thres},2} = \varrho$ which implies that $\hat{\tau}^{\text{thres},1} = \hat{\tau}^{\text{thres},2} = \tau$ and $\boldsymbol{N}_{i,t}^{\text{thres},1}(\hat{\theta}) = \boldsymbol{N}_{i,t}^{\text{thres},2}(\hat{\theta}) = \boldsymbol{N}_{i,T}(\hat{\theta})$.

Type I estimator. From the analysis of Theorem 2 (see Equation (26)), setting

$$t \geq \max\left\{\frac{\log\left(\frac{2\varrho n}{l_{n,\Theta}}\right)}{\log 2}, \frac{\log\left(\frac{(n-1)|\Theta|^2(n\Gamma_{n,\Theta}+V_{n,\Theta})}{q_1\varrho}\right)}{\log(1/a_n^\star)}\right\} = T,$$

guarantees that $\mathbb{E}\left[\boldsymbol{N}_{i,T}(\bar{\theta})\right] = \mathbb{P}\left[\boldsymbol{\nu}_{i,k,T}(\bar{\theta}) \leq \tau\right] \leq q_1$ for all $\bar{\theta} \in \bar{\Theta}$. Then following the analysis similar to Theorem 4, using the union bound and the Chernoff bound, we can prove that $\mathbb{P}[\Theta_{i,T}^{\text{thres},1} \subseteq \hat{\Theta}^\star] \leq |\bar{\Theta}|e^{-2K\pi_1^2}$. Subsequently, setting $K \geq \frac{\log(|\Theta^\star|/\alpha)}{2\pi_1^2}$ makes the error to be at most α.

Type II estimator. From the analysis of Theorem 2 (see Equation (27)), setting

$$t \geq \max\left\{\frac{\log\left(\frac{2\varrho n}{l_{n,\Theta}}\right)}{\log 2}, \frac{\log\left(\frac{(n-1)|\Theta|^2(n\Gamma_{n,\Theta}+V_{n,\Theta})}{2(1-q_2)\varrho}\right)}{\log(1/a_n^\star)}\right\} = T,$$

guarantees that $\mathbb{E}\left[\boldsymbol{N}_{i,T}(\theta^\star)\right] = \mathbb{P}\left[\boldsymbol{\nu}_{i,k,T}(\theta^\star) > \tau\right] \geq q_2$ for all $\theta^\star \in \Theta^\star$. Then following a similar analysis to Theorem 4 using the union bound and the Chernoff bound, we can prove that $\mathbb{P}[\Theta^\star \not\subseteq \hat{\Theta}_{i,T}^{\text{thres},2}] \leq |\Theta^\star|e^{-2K\pi_2^2}$. Subsequently, setting $K \geq \frac{\log(|\Theta^\star|/(1-\beta))}{2\pi_2^2}$ makes the error at most $1-\beta$.

Optimal Distributions. The proof is exactly the same as in Theorem 2.

Differentially Private Outputs. Due to the post-processing property of DP, see Dwork and Roth (2014, Proposition 2.1), the resulting estimates are ε-DP with respect to private signals.

Appendix D: Proofs and Convergence Analysis for Distributed, Private Hypothesis Testing

D.1. Proof of Proposition 1: Simple Hypothesis Testing at Significance Level α

Let $0 < \alpha', \alpha'' < 1, \alpha' + \alpha'' = \alpha$ to be determined later. The centralized UMP test at level α'' defines a threshold ϱ_c such that $\theta = 0$ is rejected if $2\Lambda(1,0) \geq \varrho_c$. The threshold ϱ_c is selected such that $\mathbb{P}[2\Lambda(1,0) \geq \varrho_c | \theta = 0] = \alpha''$. For the decentralized test, by the GM algorithm convergence analysis, we know that after K runs and T iterations given by the GM algorithm with Type I guarantee α' and threshold $\varrho^{\text{GM}} = 1$ we have the following event

$$E = \left\{\left|\frac{n}{2^{T-1}}\boldsymbol{\psi}_{i,T}(1,0) - 2\Lambda(1,0)\right| \leq \frac{n}{2^{T-1}}\right\},$$

Thus, for these values of T, K we set $\varrho_d = \varrho_c - 1$. From the above relations, we get that under $\theta = 0$ and E: $\{2\Lambda(1,0) \geq \varrho_c\} \implies \{\boldsymbol{\psi}_{i,T}(1,0) \geq \varrho_d\}$. This means that if the centralized test rejects, then the decentralized test also rejects $\theta = 0$ with probability at least $1 - \alpha'$. The Type I error is:

$$\begin{aligned}\mathbb{P}[\boldsymbol{\psi}_{i,t}^{\mathrm{GM}}(1,0) \geq \varrho_d | \theta = 0] &= \mathbb{P}[E]\mathbb{P}[\boldsymbol{\psi}_{i,t}^{\mathrm{GM}}(1,0) \geq \varrho_d | \theta = 0, E] + \mathbb{P}[E^c]\mathbb{P}[\boldsymbol{\psi}_{i,t}^{\mathrm{GM}}(1,0) \geq \varrho_d | \theta = 0, E^c] \\ &\leq \alpha'' + \alpha'.\end{aligned}$$

We choose $\alpha' = \alpha/2, \alpha'' = \alpha/2$ such that $\alpha' + \alpha'' = \alpha$. The privacy guarantee is a direct consequence of DP's post-processing property.

Q.E.D.

D.2. Extension to Composite Hypotheses

The extension to composite hypotheses is straightforward since, if we consider the generalized likelihood ratio test and run the GM algorithm with the parameters of Proposition 1 we would obtain that for T and K set as in above, with probability at least $1 - \alpha/2$ the log-belief ratio statistic

$$\left| \frac{n}{2^{T-1}} \boldsymbol{\psi}_{i,T}^{\mathrm{GM}}(\tilde{\Theta}, \tilde{\Theta}_0) - \underbrace{\sum_{i \in [n]} 2 \log \left(\frac{\sup_{\tilde{\theta}_i \in \tilde{\Theta}} \ell_i(\boldsymbol{S}_i | \tilde{\theta}_1)}{\sup_{\tilde{\theta}_0 \in \tilde{\Theta}_0} \ell_i(\boldsymbol{S}_i | \tilde{\theta}_0)} \right)}_{\text{converges to } \chi_n^2 \text{ for sufficiently large } n_i} \right| \leq \frac{n}{2^{T-1}}.$$

If $\varrho_d = F_{\chi_n^2}^{-1}(1 - \alpha/2) - 1$, then the composite hypothesis test is ε-DP and has Type I error at most α.

Appendix E: Proofs and Convergence Analysis for Distributed, Private Online Learning

E.1. Log-Belief Ratio Notations

Let $\boldsymbol{\mu}_{i,t}(\hat{\theta})$ be the beliefs for the non-private system, and let $\boldsymbol{\nu}_{i,t}(\hat{\theta})$ be the private estimates for agent $i \in [n]$ and round $t \geq 0$. For the non-private algorithm (Algorithm 5) all pairs $\hat{\theta}, \check{\theta} \in \Theta$ we let $\boldsymbol{\phi}_{i,t}(\hat{\theta}, \check{\theta}) = \log\left(\frac{\boldsymbol{\mu}_{i,t}(\hat{\theta})}{\boldsymbol{\mu}_{i,t}(\check{\theta})}\right)$ and $\boldsymbol{\lambda}_{i,t}(\hat{\theta}, \check{\theta}) = \log\left(\frac{\gamma_{i,t}(\hat{\theta})}{\gamma_{i,t}(\check{\theta})}\right)$ be the log belief ratios and the log-likelihood ratios respectively for agent $i \in [n]$ round $t \geq 0$. For the private algorithm (Algorithm 3) we let $\boldsymbol{\psi}_{i,t}(\hat{\theta}, \check{\theta}) = \log\left(\frac{\boldsymbol{\nu}_{i,t}(\hat{\theta})}{\boldsymbol{\nu}_{i,t}(\check{\theta})}\right), \boldsymbol{\zeta}_{i,t}(\hat{\theta}, \check{\theta}) = \log\left(\frac{\boldsymbol{\sigma}_{i,t}(\hat{\theta})}{\boldsymbol{\sigma}_{i,t}(\check{\theta})}\right)$, and $\boldsymbol{\kappa}_{i,t}(\hat{\theta}, \check{\theta}) = \boldsymbol{d}_{i,t}(\hat{\theta}) - \boldsymbol{d}_{i,t}(\check{\theta})$ be the private log-belief ratio, log-likelihood ratio, and noise difference ratios respectively for agent $i \in [n]$, and round $t \geq 0$. We also let the vectorized versions $\overline{\boldsymbol{\phi}}_t(\hat{\theta}, \check{\theta}), \overline{\boldsymbol{\psi}}_t(\hat{\theta}, \check{\theta}), \overline{\boldsymbol{\lambda}}_t(\hat{\theta}, \check{\theta}), \overline{\boldsymbol{\zeta}}_t(\hat{\theta}, \check{\theta})$ respectively, where the vectorization is over the agents $i \in [n]$. Finally, for each agent $i \in [n]$ and pair of states $\hat{\theta}, \check{\theta} \in \Theta$ we define the KL divergence between the states $\Lambda_i(\hat{\theta}, \check{\theta}) = \mathbb{E}_\theta\left[\log\left(\frac{\ell_i(\boldsymbol{s}_{i,0}|\hat{\theta})}{\ell_i(\boldsymbol{s}_{i,0}|\check{\theta})}\right)\right] = D_{KL}\left(\ell_i(\cdot|\hat{\theta}) | \ell_i(\cdot|\check{\theta})\right)$.

E.2. Auxiliary Lemmas

Before proving the main result for online learning, we prove the following lemma regarding the rate of convergence of the Césaro means.

LEMMA 2. *Let $\overline{\boldsymbol{X}}_0, \ldots, \overline{\boldsymbol{X}}_{t-1}$ be i.i.d. random variables (vectors), let A be an irreducible doubly stochastic matrix with the second largest eigenvalue modulus $|\lambda_2(A)|$, and let $\boldsymbol{x}_\tau = \frac{1}{n}\mathbb{1}^T\overline{\boldsymbol{X}}_\tau$ for all $0 \le \tau \le t-1$, with $\max_{1\le\tau\le t}\sqrt{\mathbb{V}\left[\boldsymbol{x}_\tau\right]} \le V$. Then*

$$\mathbb{E}\left[\left\|\frac{1}{t}\sum_{\tau=0}^{t-1}A^{t-\tau}\overline{\boldsymbol{X}}_\tau - \left(\frac{1}{t}\sum_{\tau=0}^{t-1}\boldsymbol{x}_\tau\right)\mathbb{1}\right\|_2\right] \le \frac{(n-1)V}{(1-|\lambda_2(A)|)t}.$$

Proof. For $0 \le \tau \le t-1$, we have

$$\begin{aligned}
\mathbb{E}\left[\left\|A^{t-\tau}\overline{\boldsymbol{X}}_\tau - \boldsymbol{x}_\tau\mathbb{1}\right\|_2\right] &= \mathbb{E}\left[\left\|\frac{1}{n}\mathbb{1}\mathbb{1}^T\overline{\boldsymbol{X}}_\tau + \sum_{i=2}^{n}\lambda_i(A)^{t-\tau}\overline{\boldsymbol{r}}_i\bar{\boldsymbol{l}}_i^T\overline{\boldsymbol{X}}_\tau - \boldsymbol{x}_\tau\mathbb{1}\right\|_2\right] \\
&= \mathbb{E}\left[\left\|\boldsymbol{x}_\tau\mathbb{1} + \sum_{i=2}^{n}\lambda_i(A)^{t-\tau}\overline{\boldsymbol{r}}_i\bar{\boldsymbol{l}}_i^T\overline{\boldsymbol{X}}_\tau - \boldsymbol{x}_\tau\mathbb{1}\right\|_2\right] \\
&\le \sum_{i=2}^{n}|\lambda_i(A)|^{t-\tau}\|\overline{\boldsymbol{r}}_i\|_2\|\bar{\boldsymbol{l}}_i\|_2\mathbb{E}\left[\|\overline{\boldsymbol{X}}_\tau\|_2\right] \\
&\le (n-1)|\lambda_2(A)|^{t-\tau}\mathbb{E}\left[\|\overline{\boldsymbol{X}}_\tau\|_2\right] \\
&\le (n-1)|\lambda_2(A)|^{t-\tau}\mathbb{E}\left[\|\overline{\boldsymbol{X}}_0\|_2\right] \\
&\le (n-1)|\lambda_2(A)|^{t-\tau}V.
\end{aligned}$$

Therefore,

$$\mathbb{E}\left[\left\|\frac{1}{t}\sum_{\tau=0}^{t-1}A^{t-\tau}\overline{\boldsymbol{X}}_\tau - \left(\frac{1}{t}\sum_{\tau=0}^{t-1}\boldsymbol{x}_\tau\right)\mathbb{1}\right\|_2\right] \le \frac{(n-1)V}{(1-|\lambda_2(A)|)t},$$

where we use $\sum_{\tau=0}^{t}|\lambda_2(A)|^{t-\tau} \le \sum_{\tau\ge0}|\lambda_2(A)|^{\tau} = \frac{1}{1-|\lambda_2(A)|}$. Q.E.D.

In the sequel, we show the main result for online learning.

E.3. Proof of Theorem 5: Finite-Time Guarantees for Private, Distributed Online Learning

Similarly to the asymptotic case, it suffices to pick $K = 1$. We let $\Xi_n = \mathbb{V}\left[\sum_{i=1}^{n} n_{i,\tau}\right] = \sum_{i=1}^{n}\chi_i$. We define the following sequence of "bad" events:

$$B_{n,t} = \left\{\max_{i\in[n]}\left|\frac{1}{t}\sum_{\tau=0}^{t-1}n_{i,\tau} - \xi_i\right| \ge \sqrt{\frac{\Xi_n}{\eta t}}\right\}.$$

By applying the union bound and Chebyshev's inequality, we can show that $\mathbb{P}[B_{n,t}] \le \eta$.

We condition on the "bad" event $B_{n,t}$ not happening. We can prove that, conditioned on $B_{n,t}^c$,

$$V'_{n,\Theta} = \max_{\bar{\theta}\in\bar{\Theta},\theta^\star\in\Theta^\star}\sqrt{\mathbb{V}\left[\frac{1}{nt}\sum_{i=1}^{n}\sum_{\tau=0}^{t-1}n_{i,\tau}\boldsymbol{\zeta}_{i,\tau}(\hat{\theta},\check{\theta})\middle|B_{n,t}^c\right]} \le \frac{\sqrt{2}}{n}(nQ_{n,\Theta} + V_{n,\Theta})\left(\max_{i\in[n]}\xi_i + \sqrt{\frac{\Xi_n}{\eta t}}\right).$$

The dynamics of the log-belief ratio obey $\overline{\boldsymbol{\psi}}_t(\hat{\theta},\check{\theta}) = A\overline{\boldsymbol{\psi}}_{t-1}(\hat{\theta},\check{\theta}) + \overline{\boldsymbol{\zeta}}_t(\hat{\theta},\check{\theta}) = \sum_{\tau=0}^{t-1} A^{t-\tau}\overline{\boldsymbol{\zeta}}_\tau(\hat{\theta},\check{\theta})$, where $\overline{\boldsymbol{\zeta}}_\tau(\hat{\theta},\check{\theta}) = \overline{\boldsymbol{\lambda}}_\tau(\hat{\theta},\check{\theta}) + \overline{\boldsymbol{\kappa}}_\tau(\hat{\theta},\check{\theta})$. We apply Lemma 2, and get that

$$\mathbb{E}\left[\left\|\frac{1}{t}\overline{\boldsymbol{\psi}}_t(\hat{\theta},\check{\theta}) - \frac{1}{nt}\sum_{\tau=0}^{t-1}\sum_{i=1}^{n}\boldsymbol{\zeta}_{i,\tau}(\hat{\theta},\check{\theta})\mathbb{1}\right\|_2 \middle| B_{n,t}^c\right] \leq \frac{(n-1)V'_{n,\Theta}}{t(1-|\lambda_2(A)|)}.$$

Moreover,

$$\mathbb{E}\left[\left\|\frac{1}{nt}\sum_{\tau=0}^{t-1}\sum_{i=1}^{n}\boldsymbol{\zeta}_{i,\tau}(\hat{\theta},\check{\theta}) - \frac{1}{n}\sum_{i=1}^{n}\xi_i\Lambda_i(\hat{\theta},\check{\theta})\right\|_2 \middle| B_{n,t}^c\right] \leq \frac{V'_{n,\Theta}}{t}.$$

Therefore, by Markov's inequality and the triangle inequality, we get that for every $z>0$

$$\mathbb{P}\left[\left\|\frac{1}{t}\overline{\boldsymbol{\psi}}_t(\hat{\theta},\check{\theta}) - \frac{1}{n}\sum_{i=1}^{n}\xi_i\Lambda_i(\hat{\theta},\check{\theta})\mathbb{1}\right\|_2 > z \middle| B_{n,t}^c\right] \leq \frac{V'_{n,\Theta}}{z(1-|\lambda_2(A)|)t}.$$

Therefore, we can show that by applying a union bound over $\Theta^\star$, we have that with probability $1-\eta$, for all $\bar{\theta}\in\bar{\Theta}$

$$\log \boldsymbol{\nu}_{i,t}(\bar{\theta}) \leq \boldsymbol{\psi}_{i,t}(\bar{\theta},\theta^\star) \leq -\frac{t}{n}l_{n,\Theta} + \frac{|\Theta|V'_{n,\Theta}}{2\eta(1-|\lambda_2(A)|)}.$$

To make the RHS the log-belief threshold at most some value $-\varrho$ for some $\varrho>0$, we require

$$t \geq T = \frac{\varrho + \frac{|\Theta|V'_{n,\Theta}}{2\eta(1-|\lambda_2(A)|)}}{l_{n,\Theta}/n}.$$

To determine ϱ, note that $\boldsymbol{\nu}_{i,t}(\bar{\theta}) \leq e^{-\varrho}$ for all $\bar{\theta}\neq\theta^\star$, and, thus,

$$\boldsymbol{\nu}_{i,t}(\theta^\star) = 1 - \sum_{\bar{\theta}\neq\theta^\star}\boldsymbol{\nu}_{i,t}(\bar{\theta}) \geq 1-(|\Theta|-1)e^{-\varrho}.$$

To determine a valid value of ϱ, we require $1-(|\Theta|-1)e^{-\varrho} \geq e^{-\varrho}$ which yields $\varrho \geq \log(|\Theta|)$. Therefore, setting $\varrho = \log(|\Theta|)$ and subsequently

$$T = \frac{\log(|\Theta|) + \frac{|\Theta|\frac{\sqrt{2}}{n}(nQ_{n,\Theta}+V_{n,\Theta})\left(\max_{i\in[n]}\xi_i + \sqrt{\frac{\Xi_n}{\eta}}\right)}{2\eta(1-|\lambda_2(A)|)}}{l_{n,\Theta}/n},$$

we get that

$$\mathbb{P}\left[\theta^\star = \hat{\theta}_{i,T}^{\mathrm{OL}}\right] \geq \mathbb{P}[B_{n,T}^c]\mathbb{P}\left[\theta^\star = \hat{\theta}_{i,T}^{\mathrm{OL}} | B_{n,T}^c\right] \geq (1-\eta)^2 \geq 1-2\eta.$$

Privacy. Similarly to Theorem 2, the optimal distributions are those that minimize variance subject to privacy constraints. Because there are $|\Theta|$ states, the budget ε should be divided by $|\Theta|$, resulting in $\mathcal{D}_{i,t}^\star(\varepsilon) = \mathrm{Lap}\left(\frac{\Delta_{n,\Theta}|\Theta|}{\varepsilon}\right)$. Q.E.D.

Differentially Private Outputs. Due to the post-processing property of DP (see Proposition 2.1 of Dwork and Roth (2014)), the resulting estimates are ε-DP with respect to the private signals.

Appendix F: Sensitivity of the Proportional Hazards Model

We derive the sensitivity of the proportional hazards model used in our experiments. Specifically, each center has n_i data points $\boldsymbol{S}_i$, and the treatment variable is bounded above by 1 (that is, $\max_{1\le j\le n_i}|\boldsymbol{x}_{i,j}|\le B_x=1$). The state θ obeys $|\theta|\le B_\theta$. Let $\boldsymbol{S}_i'$ be a data set of n_i points such that $\boldsymbol{S}_i$ and $\boldsymbol{S}_i'$ differ in patient k. The log partial likelihood can be written as

$$\log \ell_i(\boldsymbol{S}_i|\theta)=\log\prod_{j:\delta_{ij}=1}\frac{e^{\theta\boldsymbol{x}_{ij}}}{\sum_{r\in R(j)}e^{\theta\boldsymbol{x}_{ir}}}=\sum_{j=1}^{n_i}\boldsymbol{\delta}_{ij}\left(\theta\boldsymbol{x}_{ij}-\log\sum_{r\in R(j)}e^{\theta\boldsymbol{x}_{ir}}\right)=\sum_{j=1}^{n_i}\boldsymbol{\delta}_{ij}\theta\boldsymbol{x}_{ij}-\sum_{j=1}^{n_i}\boldsymbol{\delta}_{ij}\log\sum_{r\in R(j)}e^{\theta\boldsymbol{x}_{ir}}$$

where $R(j)=\{j':\boldsymbol{t}_{ij'}\ge\boldsymbol{t}_{ij}\}$ is the risk set of j. We are looking to bound $\max_{\boldsymbol{S}_i,\boldsymbol{S}_i':\|\boldsymbol{S}_i-\boldsymbol{S}_i'\|_1\le 1}|\log\ell_i(\boldsymbol{S}_i|\theta)-\log\ell_i(\boldsymbol{S}_i'|\theta)|$, when patient k is added, removed, or modified:

- If k is removed and k is not in any risk set $R(j)$, removing k does not affect the likelihood.
- If k is removed and k is in some risk sets, removing k affects the denominator $\sum_{r\in R(j)}e^{\theta\boldsymbol{x}_{ir}}$ for all $R(j)$ such that $k\in R(j)$.
- If k is added or modified, adding or modifying k affects the numerator $e^{\theta\boldsymbol{x}_{ik}}$ if k experiences an event, and the denominator $\sum_{r\in R(j)}e^{\theta\boldsymbol{x}_{ir}}$ if $k\in R(j)$.

Thus, the term $e^{\theta\boldsymbol{x}_{ij}}$ in the numerator is always bounded by $e^{B_\theta B_x}$ and its logarithm by $B_\theta B_x$. The denominator $\sum_{r\in R(j)}e^{\theta\boldsymbol{x}_{ir}}$ is a sum of exponentials, and its change is bounded by $e^{B_\theta B_x}$ and its logarithm by $B_\theta B_x$. Therefore, $\Delta_{n,\Theta}\le 2B_\theta B_x=2B_\theta$.

Appendix G: Lower Bounds for Distributed Hypothesis Testing

In this section, we devise information-theoretic lower bounds that are applicable to any belief-exchange algorithms. We focus on the hypothesis testing scenario with Θ consisting of a null hypothesis ($\theta=0$) and an alternative hypothesis ($\theta=1$). Initially, each agent has access to a model $\ell_i(\cdot|\theta)$ and an ε-DP privacy mechanism $\mathcal{M}_i$ which the agent uses K times to build a private signal $\boldsymbol{y}_{i,k}=\mathcal{M}_i(\boldsymbol{s}_{i,k})$. Agents can exchange the private signals $\boldsymbol{y}_{i,k}$ that they possess in each iteration T, and have access to a statistical test $\mathcal{A}:\mathcal{R}^{nK}\to\{0,1\}$ that, given private signals $\boldsymbol{y}_{1,1},\ldots,\boldsymbol{y}_{n,K}$, outputs 1 if $\theta=0$ is rejected and 0 if we fail to reject $\theta=0$. We want to find a lower bound on communication complexity, $K\cdot T$, assuming that $\mathcal{A}$ achieves a Type I error rate of α and a Type II error rate of $1-\beta$. To devise the lower bound, we rely on classic results from the bandits literature (Slivkins et al., 2019).

Theorem 6 **(Lower Bound)**. *For any belief aggregation scheme, we need at least*

$$K\cdot T\ge\frac{|1-\alpha-\beta|\mathrm{diam}\,(\mathcal{G})}{2\sum_{i\in[n]}D_{KL}\left(\hat{\ell}_i(\cdot|\theta=1)|\hat{\ell}_i(\cdot|\theta=0)\right)},\quad \textit{where}\quad \hat{\ell}_i(\boldsymbol{s}|\theta)=\ell_i(\mathcal{M}_i(\boldsymbol{s})|\theta)=\ell_i\circ\mathcal{M}_i(\boldsymbol{s}|\theta),$$

such that $\mathcal{A}$ achieves Type I error rate of α and Type II error rate of $1-\beta$.

Proof. For T, we always need $T \geq \operatorname{diam}(\mathcal{G})$ for a signal from each agent to reach any other agent. To devise a lower bound for K, let

$$\mathbb{P}\left[\mathcal{A}(\boldsymbol{y}_{1,1},\ldots,\boldsymbol{y}_{n,K})=1|\theta=0\right]=1-\alpha, \tag{31}$$

$$\mathbb{P}\left[\mathcal{A}(\boldsymbol{y}_{1,1},\ldots,\boldsymbol{y}_{n,K})=0|\theta=1\right]=\beta. \tag{32}$$

We let $P_{\theta_0}(\mathcal{B}_{i,T})=\mathbb{P}\left[\mathcal{B}_{i,T}|\theta=\theta_0\right]$ for any event $\mathcal{B}_{i,T}\subset\Omega_{i,T}$, where $\Omega_{i,T}$ is the sample space for agent i at iteration T and corresponds to all the signals observed up to time T. From simple properties of the KL divergence (see Slivkins et al. (2019)), we see that for any event $\mathcal{B}_{i,T}\subset\Omega_{i,T}$,

$$\begin{aligned}\left(P_1(\mathcal{B}_{i,T})-P_0(\mathcal{B}_{i,T})\right)^2 &\leq D_{KL}\left(P_1|P_0\right)\\ &\leq 2K\sum_{i\in[n]} D_{KL}\left(\hat{\ell}_i(\cdot|\theta=1)|\hat{\ell}_i(\cdot|\theta=0)\right),\end{aligned}$$

for $T\geq\operatorname{diam}(\mathcal{G})$, where $\hat{\ell}_i$ is the likelihood of the privatized signal, that is, $\hat{\ell}_i(\boldsymbol{s}|\theta)=\ell_i(\mathcal{M}_i(\boldsymbol{s})|\theta)$. By setting $\mathcal{B}_{i,T}=\{\mathcal{A}(\boldsymbol{y}_{1,1},\ldots,\boldsymbol{y}_{n,K})=1\}$, we get the lower bound for K.

Q.E.D.

Remark about the Randomized Response. If the mechanism corresponds to the randomized response mechanism with probability of randomization equal to $p_\varepsilon=\frac{1}{1+e^\varepsilon}$ and under $\theta=0$ the signals follow $\mathrm{Be}(1/2)$ and under $\theta=1$ the signals follow $\mathrm{Be}(1/2+\varepsilon_{\mathrm{gap}}/2)$, then we can show that on this instance, the randomized response requires

$$K\cdot T\geq\frac{2|1-\alpha-\beta|}{\varepsilon_{\mathrm{gap}}^2}\cdot\left(\frac{e^\varepsilon+1}{e^\varepsilon-1}\right)^2\cdot\frac{\operatorname{diam}(\mathcal{G})}{n},$$

because $D_{KL}\left(P_1|P_0\right)\leq\varepsilon_{\mathrm{gap}}^2(1-2p_\varepsilon)^2$, where P_0,P_1 are defined in the theorem right above. Moreover, when epsilon is small, i.e., $\varepsilon\in(0,1)$, the lower bound on the communication complexity can be approximated as $K\cdot T\geq\frac{2(1-\eta)}{\varepsilon_{\mathrm{gap}}^2}\cdot\frac{1}{\varepsilon^2}\cdot\frac{\operatorname{diam}(\mathcal{G})}{n}$, yielding an $1/\varepsilon^2$ lower bound (with respect to the privacy budget) for the randomized response mechanism.

Finding the Additive DP mechanism that Minimizes the Communication Complexity Lower Bound. For given agent models $\ell_1,\ldots,\ell_n$, finding the mechanisms that minimize the lower bound of communication complexity requires maximizing the sum of KL divergences (after applying the DP mechanism). Assuming that the noise is additive on the signal and independent of it, i.e., the private signal of each agent is $\boldsymbol{y}_i=\boldsymbol{s}_i+\boldsymbol{d}_i$, we have $\hat{\ell}_i(\cdot|\theta)=\ell_i(\cdot|\theta)*p_{\boldsymbol{d}_i}(\cdot)$ where $*$ denotes the convolution operator between the two distributions and $p_{\boldsymbol{d}_i}$ is the density of the noise. If the noise is continuous and the range of the mechanism is a subset of $\mathbb{R}^d$, the corresponding optimization problem becomes (e.g., for continuous signals and noise; see also Koufogiannis et al. (2015) and Papachristou and Rahimian (2024)):

$$\begin{aligned} \max_{p_{\boldsymbol{d}_i}(\cdot)} \quad & D_{KL}\left(\ell_i(\cdot|\theta=1) * p_{\boldsymbol{d}_i}(\cdot) | \ell_i(\cdot|\theta=0) * p_{\boldsymbol{d}_i}(\cdot)\right) \\ \text{s.t.} \quad & p_{\boldsymbol{d}_i}(\boldsymbol{u}) \geq 0, \quad \forall \boldsymbol{u} \in \mathbb{R}^d \\ & \int_{\mathbb{R}} p_{\boldsymbol{d}_i}(\boldsymbol{u}) d\boldsymbol{u} = 1 \\ & \|\nabla_{\boldsymbol{s}} \log p_{\boldsymbol{d}_i}(\boldsymbol{u}-\boldsymbol{s})\|_\infty \leq \varepsilon, \quad \forall \boldsymbol{u} \in \mathbb{R}^d \end{aligned} \tag{33}$$

where $\ell_i(\cdot|\theta) * p_{\boldsymbol{d}_i})(\boldsymbol{s}) = \int_{\mathbb{R}^d} p_d(\boldsymbol{u}) \ell_i(\boldsymbol{s}-\boldsymbol{u}|\theta) d\boldsymbol{u}$.

Appendix H: Additional Simulation Experiments

H.1. Toy Example: MLE and OL for a Single Treatment

We test the MLE and OL algorithms in an arrangement of $n = 5$ hospitals possible values $\Theta = \{0, -\log 2\}$, where the value of $-\log 2$ examines whether the patients are 1/2 less likely to die. We use a privacy budget of $\varepsilon = 1$ and error bounds equal to $\alpha = 1 - \beta = 0.05$. The evolution of the beliefs and the beliefs at the terminal time T – calculated by applying Theorem 2 – are shown in Figure 1(a). Algorithm 1 is able to successfully recover $\Theta^\star = \{-\log 2\}$ which is the true maximizer, that is, $\Lambda(-\log 2) > \Lambda(0)$. The DP algorithms are compared against the centralized baseline, which corresponds to all centers sharing their signals without privacy. In addition, Figure 1(b) shows the result of applying the two threshold algorithm (with a single threshold), which is also able to recover successfully $\Theta^\star$, similar to the AM/GM algorithm.

We perform a similar analysis in an online manner, where we assume that the centers perform survival analysis every 10 days (we achieve that by splitting the data evenly among these days). In Figure 1(c), we report the time-averaged log-belief ratio as well as the beliefs at the terminal time T for the same choice of topology and privacy budget. The time-averaged log-belief ratio is shown against the limiting value of the non-DP baseline (see Rahimian and Jadbabaie (2016b)). Again, agents collectively identify the true state $\theta^\circ = -\log 2$.

H.2. Time and Space Complexity

We first start by analyzing the time complexity of our method. Specifically, in the distributed MLE setting, if agent i has n_i data points and model ℓ_i which can be computed in $\mathcal{T}_i(n_i)$ for one value of $\theta \in \Theta$, then the distributed MLE is parallelizable between the agent and the total time complexity per agent is $\mathcal{O}\left(|\Theta| K \left(\mathcal{T}_i(n_i) + T \deg_{\mathcal{G}}(i)\right)\right)$ which includes an initialization cost of $|\Theta| K \mathcal{T}_i(n_i)$ to calculate noisy likelihoods. For example, in the case of the distributed proportional hazards model, the likelihoods can be computed in $\mathcal{T}_i(n_i) = \mathcal{O}\left(n_i \log n_i\right)$. Then, there is a communication cost, which consists of the number of belief exchanges dictated by the communication complexity KT, and the cost of propagation for each iteration, which is $\mathcal{O}(|\Theta| K \deg_{\mathcal{G}}(i))$ for all states. Similarly, in the OL

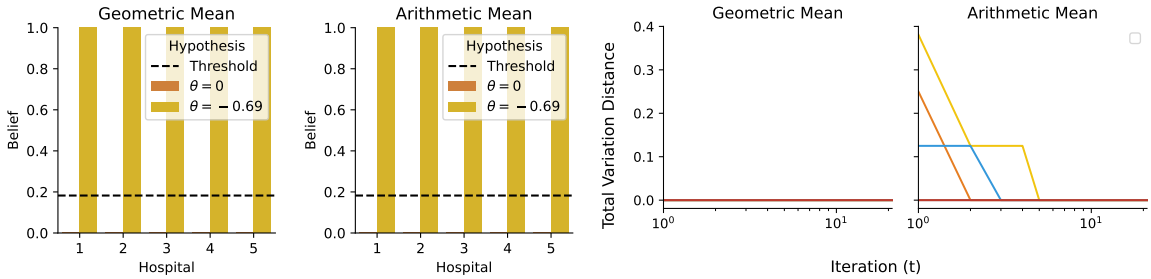

(a) AM/GM Algorithm

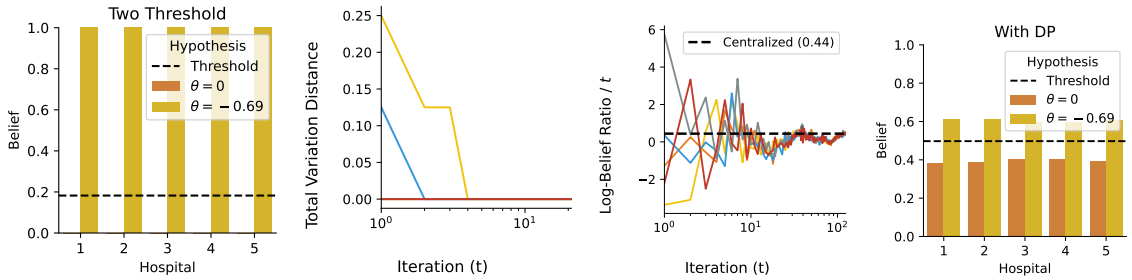

(b) Two Threshold Algorithm (c) Online Learning

Figure H.1 **Top (a):** Resulting beliefs and total variation distance between the beliefs and the ground truth for the GM and AM estimators (Algorithm 1) for the ACTG study data for $n = 5$ centers examining the effect the ddI treatment on patient survival assuming a proportional hazards model. Section 2.1 describes the model. The topology between the hospitals is taken to be the complete graph. We have set $\varepsilon = 1$, $\alpha = 1 - \beta = 0.05$ and $\varrho^{\mathrm{AM}} = \varrho^{\mathrm{GM}} = 1.5$. The resulting estimators recover $\Theta^\star = \{-\log 2\}$. The number of iterations T and the number of rounds K have been computed according to Theorem 2.

Bottom Left (b): Resulting beliefs and total variation distance between the beliefs and the ground truth for the Two Threshold algorithm. We have set $\varepsilon = 1$, $\alpha = 1 - \beta = 0.05$ and $\hat{\tau}^{\mathrm{thres},1} = \hat{\tau}^{\mathrm{thres},2} = 0$ and $\tau^{\mathrm{thres},1} = \tau^{\mathrm{thres},2} = 1.5$ (single-threshold recovery; cf Corollary 1). The resulting estimators recover $\Theta^\star$ successfully. The number of iterations T and number of rounds K have been computed according to Theorem 4.

Bottom Right (c): Resulting log-belief ratios and terminal beliefs for the online learning algorithm on the ACTG study. We assume that the centers exchange beliefs on a daily basis. The dashed line corresponds to the value that the time-averaged log-belief ratio converges in the non-DP regime.

regime, the total time complexity per agent is $\mathcal{O}\left(|\Theta| \sum_{t=1}^{T} (\mathcal{T}_i(n_{i,t}) + \deg_{\mathcal{G}}(i))\right)$. Finally, in terms of space complexity, each agent can simply maintain their belief in each state, giving a space complexity $\mathcal{O}(K|\Theta|)$ during communication.

H.3. Runtime Comparison with Homomorphic Encryption Methods

To further test the practical applicability of our method, we compare the runtime of our method with existing holomorphic encryption methods (HE) and, in particular, the FAMHE method introduced in Froelicher et al. (2021) as the number of providers (n) grows and the number of data grows. Both our

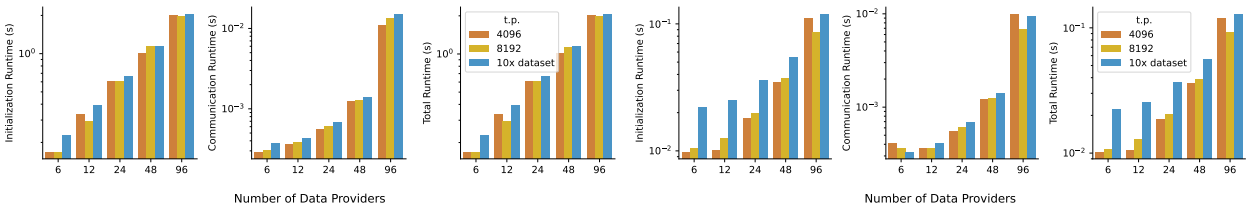

(a) Cancer tumor mutational load study (b) ACTG study 175

Figure H.2 Runtime of distributed MLE algorithm for (a) the study by Samstein et al. (2019) on the left and (b) the ACTG study on the right, for varying values of n used in measuring the performance of the FAMHE method Froelicher et al. (2021). The privacy budget is set to $\varepsilon = 1$ (tight privacy), and the error rates are set to $\alpha = 1 - \beta = 0.05$, and the thresholds are set to $\varrho^{\text{AM}} = \varrho^{\text{GM}} = 0.1$. In accordance with Froelicher et al. (2021), we have constructed 3 datasets: a dataset consisting of 4096 timepoints (t.p.) via resampling the original data with replacement, a dataset of 8192 t.p. via resampling the original data with replacement, and a dataset which consists of 10 times the original data.

method and FAMHE are capable of performing survival analysis in a distributed regime; however, they apply different privacy protections and cannot be compared prima facie. Generally, HE methods have the strongest possible privacy. However, these protections fall short in scalability considerations, for which DP with a small ε is a better alternative in terms of efficiency. Furthermore, the most recent HE methods designed for multicenter trials do not have open-source implementations, making direct comparison difficult (see, e.g. Froelicher et al. (2021) and Geva et al. (2023)).

First, our aim is to compare the two methods in the same dataset and experiment, which corresponds to distributed survival analysis with cancer data introduced in Samstein et al. (2019). Specifically, Samstein et al. (2019) finds that tumor mutational burden (TMB) is a significant predictor of survival outcomes in patients with metastatic cancer undergoing immune checkpoint inhibitor (ICI) therapy. Analyzing data from 1,662 advanced cancer patients, the authors demonstrated that higher TMB levels are associated with better survival prospects.

Froelicher et al. (2021) use the data from Samstein et al. (2019) and perform survival analysis (see Figure 2 in their paper) by splitting the data equally among n providers. Then, the authors report the runtime as a function of n for three datasets: with 4096 time points (t.p.), 8192 t.p., and a dataset that consists of 10 times the original data, and report the total runtime in Figure 2(c). To construct the fairest possible comparison, we consider the same dataset and the distributed proportional hazards model with TMB as a covariate and two hypotheses – null ($\theta = 0$) and a composite alternative ($\theta \neq 0$) – similarly to our initial example for ACTG (see Section 2.1). The first two datasets (4096 and 8192 t.p.) are constructed by resampling the original data with replacement. We choose ε to be 0.1, which corresponds to a very strong privacy regime, significantly less than the privacy budgets used in the 2020 US Census Bureau (2021) and other healthcare contexts (Ficek et al., 2021; Dwork et al., 2019).

Finally, we chose the error bounds to be no more than 0.1. In Figure H.2, we report the per-agent cost of inference for $n \in \{6, 12, 24, 48, 96\}$ in a fully connected topology, analogous to Figure 2(c) of Froelicher et al. (2021). Our reported runtimes range from $\sim 10^{-2}$s to ~ 1s, which is a $10\times - 1000\times$ improvement over the runtimes reported in Froelicher et al. (2021).

H.4. Comparison with First-order Methods

We compare with the first-order method of Rizk et al. (2023) with graph-homomorphic noise. The distributed optimization problem we are solving is

$$\max_{\theta \in \mathbb{R}} \frac{1}{n} \sum_{i \in [n]} 2 \left(\log \ell_i(\boldsymbol{S}_i | \theta) - \log \ell_i(\boldsymbol{S}_i | 0) \right).$$

The distributed updates on the parameters with graph-homomorphic noise according to Rizk et al. (2023) are equivalent to

$$\boldsymbol{\theta}_{i,t} = \sum_{j \in [n]} a_{ij} \boldsymbol{\theta}_{j,t-1} + \eta_{\text{lr}} \text{Clip} \left(\frac{\mathrm{d} \log \ell_i(\boldsymbol{S}_i | \boldsymbol{\theta}_{i,t-1})}{\mathrm{d}\boldsymbol{\theta}_{i,t-1}}, 2B_\theta B_x \right) + \text{Lap} \left(\frac{2B_x B_\theta T \eta_{\text{lr}}}{\varepsilon} \right),$$

where η_{lr} is the learning ratem and the Clip operation clips the gradient to have L_1 norm at most $2B_\theta B_x$. To derive the Laplace noise, note that the sensitivity of the clipped gradient is, at most, $2\eta_{\text{lr}} B_\theta B_x$. Additionally, since the per-agent budget is ε, we need to scale the budget by T. In the simulations, we use $B_\theta = 1$, $\eta_{\text{lr}} = 0.001$, and in order to have a fair comparison to our algorithm, we set T to be equal to the number of iterations we run our algorithms. For large sample sizes the P values are given as $P_i = 1 - F_{\chi^2_1} \left(2 \left[\log \ell_i(\boldsymbol{S}_i | \boldsymbol{\theta}_{i,T}) - \log \ell_i(\boldsymbol{S}_i | 0) \right] \right)$.

H.5. Additional Experimental Results with the Cancer Dataset

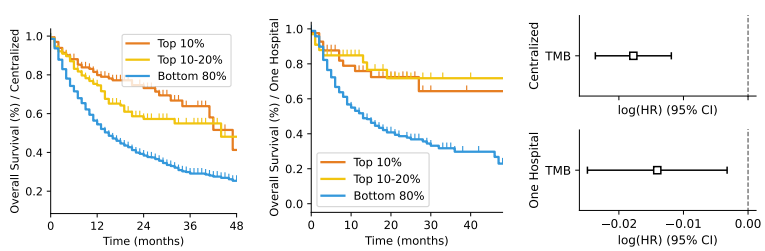

(a) Centralized curves (b) One hospital (c) log Hazard Ratios

Figure H.3 **Left:** Kaplan-Meier survival curves for the cancer data. The three curves correspond to high tumor mutational burden/TMB (top 10%), medium TMB (top 10%-20%) and low TMB (bottom 80%). **Middle:** Survival curves for one hospital (the data is split equally among 5 hospitals) for the same study. **Right:** Log hazard ratios with 95% confidence intervals from the fitted proportional hazards model for centralized and one hospital.

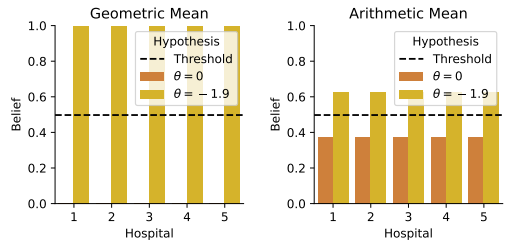

Figure H.4 Resulting beliefs $\boldsymbol{\nu}_{i,T}^{\mathrm{GM}}(\theta)$ and $\boldsymbol{\nu}_{i,T}^{\mathrm{GM}}(\theta)$ for the GM and AM estimators (Algorithm 1) for the cancer study data for $n = 5$ fully connected centers examining the effect of TMB on patient survival assuming a proportional hazards model and $\Theta = \{0, \log(0.15)\}$. We have set $\varepsilon = 1$, $\alpha = 1 - \beta = 0.05$, and $\varrho^{\mathrm{AM}} = \varrho^{\mathrm{GM}} = 0.1$. The resulting estimators yield $\{\log(0.15)\}$ as the MLE set which agrees with the ground truth.

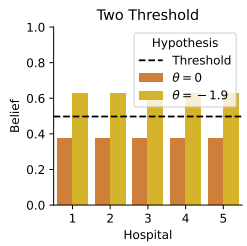

Figure H.5 Resulting beliefs $\boldsymbol{N}_{i,T}(\theta)$ for the two threshold algorithm (Algorithm 2) for the cancer study data for $n = 5$ fully connected centers examining the effect of TMB on patient survival assuming a proportional hazards model. We have set $\varepsilon = 1$, $\alpha = 1 - \beta = 0.05$ and $\hat{\tau}^{\mathrm{thres},1} = \hat{\tau}^{\mathrm{thres},2} = 0$ and $\tau^{\mathrm{thres},1} = \tau^{\mathrm{thres},2} = 0.5$ (single-threshold recovery; cf Corollary 1). The resulting estimators yield $\{\log(0.15)\}$ as the MLE set. The number of iterations T and number of rounds K have been computed according to Theorem 4.

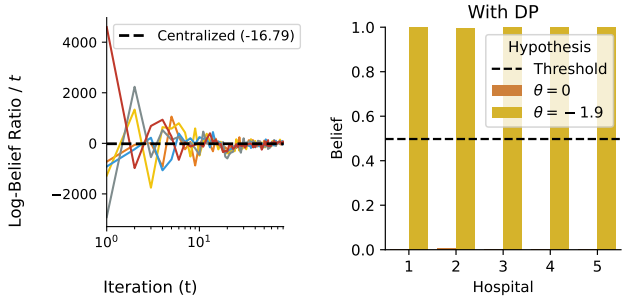

Figure H.6 Resulting log-belief ratios and terminal beliefs for the online learning algorithm on the cancer study. We assume that the centers exchange beliefs on a daily basis for $\varepsilon = 1$. Data is evenly split among $n = 5$ centers. The dashed line corresponds to the value that the time-averaged log-belief ratio converges in the non-DP regime.

Appendix I: Table of Notations

x	scalar
$\overline{\boldsymbol{x}}$	vector
$\boldsymbol{x}$	element of vector $\overline{\boldsymbol{x}}$
$\|\overline{\boldsymbol{x}}\|_p$	ℓ_p-norm of $\overline{\boldsymbol{x}}$
X	matrix
n	number of agents
ε	privacy budget
α, $1-\beta$	Type I, Type II error probabilities
η	$\eta = \max\{\alpha, 1-\beta\}$
A	communication matrix
Θ	set of states
$\Theta^\star$, $\bar{\Theta}$	set of MLE, non-MLE states
$\theta^\star$, $\bar{\theta}$, θ°	MLE, non-MLE/false state, true state (in the online learning case)
$f^\star$	density of MLE states ($f^\star = \|\Theta^\star\|/\|\Theta\|$)
$\ell_i(\boldsymbol{s}\|\hat{\theta})$	likelihood of agent i given signal $\boldsymbol{s}$ at state $\hat{\theta}$
T, K	number of iterations, number of rounds
$\boldsymbol{\gamma}_{i,t}(\hat{\theta})$	joint likelihood of agent i given data $\boldsymbol{S}_{i,t}$ of size $n_{i,t}$ at state $\hat{\theta}$ in the non-private regime
$\boldsymbol{\lambda}_{i,t}(\hat{\theta}, \check{\theta})$	joint log likelihood ratio of agent i at iteration t between states $\hat{\theta}$ and $\check{\theta}$ in the non-private regime
$\Lambda(\hat{\theta}, \check{\theta})$	joint log likelihood ratio between states $\hat{\theta}$ and $\check{\theta}$ over all agents and data in the non-private regime
$\boldsymbol{\sigma}_{i,t}(\hat{\theta})$	joint likelihood of agent i given data $\boldsymbol{S}_{i,t}$ of size $n_{i,t}$ at state $\hat{\theta}$ in the private regime
$\boldsymbol{\zeta}_{i,k,t}(\hat{\theta}, \check{\theta})$	joint log likelihood ratio of agent i at iteration t at round k between states $\hat{\theta}$ and $\check{\theta}$ in the private regime
$Z_k(\hat{\theta}, \check{\theta})$	joint log likelihood ratio between states $\hat{\theta}$ and $\check{\theta}$ at round k over all agents and data in the private regime
$\boldsymbol{\mu}_{i,t}(\hat{\theta})$	belief of agent i at iteration t on state $\hat{\theta}$ in the non-private regime
$\boldsymbol{\nu}_{i,k,t}(\hat{\theta})$	belief of agent i at iteration t at round k on state $\hat{\theta}$ in the private regime
$\boldsymbol{\phi}_{i,t}(\hat{\theta}, \check{\theta})$	log belief ratio of agent i at iteration t on pair of states $\hat{\theta}$ and $\check{\theta}$ in the non-private regime
$\boldsymbol{\psi}_{i,k.t}(\hat{\theta}, \check{\theta})$	log belief ratio of agent i at iteration t at round k on pair of states $\hat{\theta}$ and $\check{\theta}$ in the private regime
$\boldsymbol{d}_{i,t}(\hat{\theta})$, $\boldsymbol{d}_{i,k}(\hat{\theta})$	noise variable for agent i, state $\hat{\theta}$, and iteration t, or round k
$\mathcal{D}_{i,t}(\hat{\theta};\varepsilon)$, $\mathcal{D}_{i,k}(\hat{\theta};\varepsilon)$	noise distribution for agent i, state $\hat{\theta}$, and iteration t, or round k
$n_{i,\tau}$	random number of signals that i observes at time step τ in the online learning setup
ξ_i	expected number of signals that i observes at each time step in the online learning setup: $\xi_i = \mathbb{E}[n_{i,\tau}]$
χ_i	variance of number of signals that i observes at each time step in the online learning setup: $\chi_i = \mathbb{V}[n_{i,\tau}]$
$\boldsymbol{\delta}_{ij}$	censoring indicator for patient j in center i
$\boldsymbol{x}_{ij}$	treatment indicator for patient j in center i
$\boldsymbol{t}_{ij}$	survival or censoring time in days for patient j in center i
$b_n^\star = \|\lambda_2(A)\|$	second-largest eigenvalue modulus of A
$a_n^\star$	second largest eigenvalue modulus of $(A+I)/2$, also denoted by $\|\lambda_2(A+I)\|/2$
$\Gamma_{n,\Theta}$	largest likelihood magnitude
$\Delta_{n,\Theta}$	global sensitivity
$l_{n,\Theta}$	minimum KL-divergence between non-MLE and MLE states (false and true states in the OL case)
Ξ_n	sum of variances of numbers of signals across all agents
$V_{n,\Theta}$	sum of standard deviations of the noises over all agents
$Q_{n,\Theta}$	maximum standard deviation of KL-divergence for any agent and any pair of states