Optimal Data Acquisition
with Privacy-Aware Agents

Abstract—We study the problem faced by a data analyst or platform that wishes to collect private data from privacy-aware agents. To incentivize participation, in exchange for this data, the platform provides a service to the agents in the form of a statistic computed using all agents' submitted data. The agents decide whether to join the platform (and truthfully reveal their data) or not participate by considering both the privacy costs of joining and the benefit they get from obtaining the statistic. The platform must ensure the statistic is computed differentially privately and chooses a central level of noise to add to the computation, but can also induce personalized privacy levels (or costs) by giving different weights to different agents in the computation as a function of their heterogeneous privacy preferences (which are known to the platform). We assume the platform aims to optimize the accuracy of the statistic, and must pick the privacy level of each agent to trade-off between i) incentivizing more participation and ii) adding less noise to the estimate.

We provide a semi-closed form characterization of the optimal choice of agent weights for the platform in two variants of our model. In both of these models, we identify a common nontrivial structure in the platform's optimal solution: an instance-specific number of agents with the least stringent privacy requirements are pooled together and given the same weight, while the weights of the remaining agents decrease as a function of the strength of their privacy requirement. We also provide algorithmic results on how to find the optimal value of the noise parameter used by the platform and of the weights given to the agents.

Index Terms—online platforms, data acquisition, privacy-aware agents, differential privacy, endogenous participation.

I. INTRODUCTION

Recent advancements in machine learning algorithms and large-scale computation has reaffirmed the crucial value of information, leading to unprecedented levels of data gathering. For example, the recommendation systems used by platforms like Netflix, TikTok, or YouTube are trained on massive amounts of data regarding user behavior and preferences. However, this accumulation of information has raised important concerns regarding the privacy costs suffered by the users that this information pertains to. To mitigate this issue, a large body of research has focused on designing algorithms that process the sensitive user information while limiting their incurred privacy costs (a prominent line of such work focuses on *differential privacy*). The main limitation of this approach is that the reduced privacy costs often come at the expense of lower quality outcomes (e.g., a recommendation system with very strong privacy guarantees may yield poor recommendations), hurting the same users that it is aiming to protect. Our goal in this paper is to develop a better understanding of the trade-offs that such users face between privacy costs and the resulting quality of service, and to design optimal data acquisition mechanisms which respect the users' preferences.

The privacy cost that agents may suffer by releasing access to their data is well-studied: e.g., users in online platforms and social media applications may not want to reveal their search and watch histories or content preferences. In response to these concerns, many platforms allow their users to opt out of sharing their data (e.g., a YouTube user can opt out of letting the platform track their activity). Another common example where collecting sensitive information may be really valuable are medical studies aimed at developing a better understanding of some rare disease. Clearly, an individual that has this disease may be reluctant to share this information, and prefer not to participate in the study.

What may not be as well-understood in these examples is the impact of the potential non-monetary benefits that the agents can accrue by contributing their data: e.g., a prime motivation to participate in a medical study about a rare disease is the hope that it may lead to new treatments, which would directly benefit those who suffer from it. Similarly, by revealing their content preferences, users of platforms like Netflix or Youtube can help these platforms improve their recommendation engines which, in turn, provides these users with a higher quality of service. In general, the more significant the potential benefits are, the less reluctant the agents are to share their data. Furthermore, it is often the case that these benefits increase (and privacy costs drop) when more users participate, giving rise to interesting and, to the best

of our knowledge, less well-understood complementarity phenomena across users.

Our goal in this paper is to model and analyze such settings in which agents can decide whether or not to release access to their data, by considering both the benefit they would obtain and the privacy losses they would incur. Specifically, we approach this problem from the perspective of a platform whose goal is to maximize the value of the final computation (e.g., the quality of service of a system or the accuracy of a study), while respecting the preferences of the agents. To achieve this goal, the platform can determine the extent to which it will introduce differential privacy protections, taking into consideration the agents' preferences and aiming to incentivize their participation.

A. Summary of contributions

In this paper, for simplicity, we consider a learner or platform that is interested in performing a simple estimation task: understanding the mean of a population distribution. We see this simple estimation task as a possible proxy for more complex machine learning tasks (such as training a recommendation system), and leave the study of such machine learning tasks to future work. We consider a setting in which the learner controls two types of variables: i) the amount of noise η centrally added to the computation for differential privacy, and ii) the weights $w_1, \ldots, w_n$ given to the data of agents $1, \ldots, n$ in the learner's computation. By giving different weights to different agents, the platform can provide personalized privacy levels to agents with varying privacy attitudes; see the preliminary Section III for more details on how the privacy level obtained by an agent i depends on η and w_i. We are interested in designing the weights $\vec{w}$ and the noise η to optimize the accuracy of the learner's statistic.

Ultimately, the accuracy of this statistic will depend on the participation decisions of the agents: as more agents participate, the learner is able to collect more data and to refine his statistic. The first main contribution of our paper is to propose two potential models of how agents decide whether to participate in the platform:

- In Section IV-A, we introduce the "quasi-linear" agent model. In this model, an agent explicitly trades-off the privacy losses they incur with the benefit they get from the platform. They only decide to participate in the platform if the anticipated benefit is higher than the cost for sharing their data. A version of this model is also used in the concurrent work of [1].

- In Section IV-B, we introduce a simpler variant of our model, called the "privacy-constrained" model. In this model, an agent is willing to join the platform as long as i) they get some benefit from it and ii) a minimum privacy requirement (that may be different for different agents) is met.

We then proceed to characterizing the optimal choice of estimator (i.e. of weights $\vec{w}$ and noise parameter η) in semi closed-form:

- In Section V, we do so for the quasi-linear agent participation model. There, we remark that the optimal solution has a non-trivial structure, similar to that of [2]: namely, the agents with the least stringent privacy requirements are pooled together and given the same weight, while agents with higher privacy requirements are given weights that decrease with the strength of their privacy attitudes. We also provide algorithmic guidance on how to find the optimal value of η. We note that how to elicit the agents' privacy costs when they are strategic and can misreport these costs is discussed in [1].

- In Section VI, we show that a similar structure (a pooling region followed by a decreasing weight as privacy attitudes become more stringent) arises in the alternative, privacy-constrained model. We provide expressions for both $w_1, \ldots, w_n$ and η nearly in closed-form, up to a single unknown parameter t which controls the number of agents that are pooled together. We also remark that this variant of the model has simple incentive properties: it is in the agents' best interest to report their privacy costs truthfully, even without interventions or payments by the learner.

II. RELATED WORK

Recently, there has been a lot of interest in the study of data transactions in the computer science, operations research, and economics literatures. For example, [3], [4] study how to model and design data markets.

Much of the literature focuses on one major building block for data transactions–deciding how to efficiently and optimally acquire data from a collection of agents (or "data providers"). On main focus of this literature is settings in which the data providers must be compensated for their data. For example, [5], [2], [6], [7], [8], [9], [10], [11] look at the pricing and purchase of such data when the provided data is verifiable (but providers may be strategic and lie about their costs for revealing their data). There is also a significant line of work–such as [12],

[13], [14], [15], [16], [17], [18], [19], [20]–on the case of non-verifiable data points, where providers can also lie about their data in order to steer the learner's model towards desired outcomes.

A significant part of this literature singles out privacy loss as the main reason data providers must be compensated for their data. This gave rise to a body of work that focuses on data acquisition under differential privacy constraints, e.g. [21], [22], [23], [24], [25], [26], [27], [28]. I.e., the seller must provide formal privacy guarantees on how the providers' data is used, while often still compensating them for any remaining privacy losses. This is where our work lies; we adopt the same point of view as [22], [28] in that we consider settings in which agents have an inherent interest in the statistic or service offered by the platform that is trained on their data, rather than solely in the payments they receive from the platform.

One of the main, salient elements of our model is that the quality of the estimate or service provided by the platform depends not only on the privacy level that the platform offers, but also on the number of providers that join the platform and report their data. In turn, the agents' participation decisions are an endogenous aspect of our model, as in the works of [29] and [1]. Similarly to our setting, [29] consider a setting in which the privacy cost a data provider incurs depends on other providers' participation decision; the main distinction compared to our work is that in [29], agents only care about how much privacy they obtain, not on how the collected data is used by the buyer or platform to offer a useful service or statistic in return. The new work of [1] considers a data acquisition setting with verifiable data where agents obtain a benefit that depends on the accuracy of the platform's model, rather than only from payments they get from their data. [1]'s model is similar and contemporaneous to ours; while we assume that the privacy preferences of the agents were known, [1] considers settings where agents strategically report and may lie about their privacy costs. We remark that our results are mostly orthogonal to theirs: they provide novel algorithms to solve the mechanism design problem of incentivizing truthful cost reporting while optimizing the accuracy of the platform's estimate; we focus on understanding properties of and on characterizing how much the optimal estimator for the platform weights each agent's data in semi- closed form, as a function of their privacy preferences. We also incorporate several additional modeling elements relative to [1]: in particular, i) we assume that each agent may benefit from the platform's estimation in a possibly non-linear way, and ii) we provide a second, alternative model of agents' privacy preferences and of how they decide to participate in the platform.

III. DIFFERENTIAL PRIVACY PRELIMINARIES

In this paper, we focus on *differential privacy* as our main privacy technique. Differential privacy was first introduced in the seminal work of [30] and aims to prevent an attacker from being able to infer an agent's data by observing or post-processing the output of an algorithm, e.g. the output of a learner's statistical computation or machine learning model. In this section, we focus on presenting the minimal knowledge of differential privacy needed for this paper; for a more detailed discussion of differential privacy, please refer to [31].

Differential privacy protects an agent's data by comparing two possible worlds for each agent; the difference between these two worlds is that they consider two possible different values for the data of this agent. Differential privacy requires that one (almost) cannot distinguish between these two worlds by looking at the (distribution over) outputs of the learner's computation; i.e., one cannot tell with any reasonable certainty what the data point of the agent was, since the outcome of the computation (nearly) does not depend on its value. Formally, a learner runs a computation of mechanism $\mathcal{M}$ which takes a dataset x as an input, and outputs some function or property $\mathcal{M}(x)$ of that dataset. Given n agents whose data is used in the learner's mechanism, one can think of a dataset x as a vector of entries $(x_1, \ldots, x_n)$, where x_i is the data of agent i. We first introduce the definition of neighboring datasets:

Definition 1. *Two datasets x and x' are neighboring with respect to agent i (or "i-neighbors") if they differ only in agent i's data. I.e., $x_j = x'_j$ for all $j \neq i$.*

Differential privacy, as informally described above, requires that the outputs of mechanism $\mathcal{M}$ differ little on any two neighboring databases x and x'. This is formalized as follows:

Definition 2 (ε-differential privacy)**.** *Let $\varepsilon > 0$. A randomized algorithm $\mathcal{M}$ is ε-differentially private with respect to agent i if for any outcome set $O \subset Range(\mathcal{M})$ and for all neighboring databases x, x' with respect to i,*

$$\Pr[\mathcal{M}(x) \in O] \leq \exp(\varepsilon) \Pr[\mathcal{M}(x') \in O].$$

Here, the parameter ε controls how much privacy each agent gets. As ε decreases, $\exp(\varepsilon)$ also decreases

and the above constraint becomes more and more stringent, improving the level of privacy guaranteed by the mechanism. For $\varepsilon = 0$, it in fact requires that $\Pr\left[\mathcal{M}(x) = o\right] = \Pr\left[\mathcal{M}(x') = o\right]$; i.e., the outcome of the mechanism is independent of the input data and thus perfectly preserves privacy. As $\varepsilon \to +\infty$, the above constraint is trivially satisfied by any mechanism and no privacy protection is provided.

One of the simplest way to answer a desired statistical query in a differentially private manner is to add noise to the output of said query. Intuitively, as the amount of added noise *increases*, the dependency of the result of said query on any particular agent's data *decreases* (equivalently, the level of privacy obtained by agents *increases*). The most basic and common mechanism to obtain differential privacy answers to numerical queries is the Laplace mechanism, which adds Laplace noise to the output of a query.

Definition 3. *Let q be a numerical query, i.e. $q(x) \in \mathbb{R}$ for all x. The Laplace mechanism is defined as*

$$\mathcal{M}_L(x, q, \eta) = q(x) + Z,$$

where Z is a random variable drawn from the Laplace distribution with parameter η.

The level of privacy obtained by the Laplace mechanism depends on the sensitivity of the query we aim to answer; i.e., how much the value of this query changes when the data entry of a single agent in the database changes. Formally, the sensitivity of a query with respect to agent i is defined as

$$(\Delta q)_i = \max_{x, x' \text{ i-neighbors}} |q(x) - q(x')|.$$

We then have the following privacy guarantee for agent i:

Definition 4. $\mathcal{M}_L(x, q, \varepsilon) = q(x) + Z$ *is $\eta (\Delta q)_i$-differentially private with respect to agent i.*

Finally, we note that our goal is to both provide individual agents reporting their data with privacy guarantees while at the same time obtaining an accurate estimate of the statistic we are interested in. Because we consider unbiased estimators in this paper, the accuracy of said estimator is directly linked to its variance. The variance of the Laplace mechanism with parameter η on query q is given by

$$\operatorname{Var}\left(q(x) + Z\right) = \operatorname{Var}_x \left(q(x)\right) + \frac{2}{\eta^2}.$$

IV. Model

We model a setting in which a data analyst or platform aims to incentivize privacy-aware agents to share their data with or join the platform, then collects their data and uses it to estimate a statistic. To incentivize agent participation, the platform simultaneously aims to provide privacy guarantees to agents who join the platform while also offering a useful service to the agents who join through their machine learning model. E.g., the machine learning model could be a platform's recommendation system, such as the ones offered by platforms such as YouTube and TikTok; it could also be the product of a medical study on a rare disease, where the individuals contribute their sensitive medical data to a study in the hopes of getting better treatments and medical outcomes in return.

The platform faces a population of n agents. Each agent has a private data point d_i. The data points are drawn i.i.d. from an unknown distribution with unknown mean μ but known variance σ^2. Each agent also has a linear privacy cost function given by $c_i \varepsilon_i$, where $c_i \geq 0$ is an agent-specific scalar and $\varepsilon_i > 0$ is the level of differential privacy obtain by agent i if he joins the platform; this linearity assumption follows that of [25].

The goal of the platform is to i) incentivize agents to join the platform, then ii) compute an unbiased estimator $\hat{\mu}$ of μ. The platform wants this estimator to be as accurate as possible. Letting $S \in [n]$ be the set of agents that decide to join the platform, we assume that the platform's estimator is linear, i.e. given by

$$\hat{\mu}(S, \mathbf{w}, \eta) = \sum_{i \in S} w_i d_i + Z(\eta),$$

where w_i is the weight assigned to the data of agent i and Z is a random variable drawn from a Laplace distribution with parameter $\eta \geq 0$ for privacy. We denote as $\mathbf{w}$ the vector of all w_i's. Since we require our estimator to be unbiased, we assume that $w_i \geq 0$ for all $i \in [n]$ and that $\sum_{i \in S} w_i = 1$. The platform optimizes over both the choice of weights $\{w_i\}_{i \in S}$ and of noise parameter η.

Because the estimator used by the platform is unbiased, we can measure its performance (here, its expected mean-squared error) through its variance. The variance of $\hat{\mu}$, as per preliminary section III, is given by

$$\operatorname{Var}(\hat{\mu}) = \sum_{i \in S} w_i^2 \sigma^2 + \frac{2}{\eta^2}.$$

The order of operations is then the following:

4

1) The analyst announces the weight vector $\mathbf{w}$ and the noise parameter η that she will use in the computation.
2) Each agent i decides whether he wants to participate given $\mathbf{w}, \eta$.
3) The analyst computes the estimator $\hat{\mu}$ on the participating agents.

We propose two variants on how we model agents' privacy attitudes, utilities, and participation decisions. In the "quasi-linear agent" model, agents maximize a quasi-linear utility functions that trades-off the quality of the final model and their privacy costs. In the "privacy-constrained agent" model, agents aim to maximize the utility they get from the platform's model under a constraint that their privacy is not violated by more than a desired tolerance.

A. The Quasi-Linear Agent Model

In the quasi-linear model, agent i has a quasi linear utility for participating in the platform, which trades-off his privacy cost for reporting his data and his utility from the platform's estimation. Noting that the sensitivity of estimator $\hat{\mu}$ with respect to agent i is given by w_i, the level of privacy obtained by agent i is given by $\varepsilon_i = w_i \eta$ (as discussed in Section III), and i incurs cost $c_i w_i \eta$ for participating in the platform. In turn, we consider the following quasi-linear utility for the agent:

$$u_i(\mathbf{w}, \eta) = f\left(\sigma^2 \sum_{i \in S} w_i^2 + \frac{2}{\eta^2}\right) - c_i w_i \eta$$

for some decreasing function f; i.e., agent i's utility increases when the variance of the platform's model decreases, and when his privacy cost decreases. If the agent does not join the platform, we assume they have access to an outside option with utility o (for example, they could use their own data point as an estimate). The agent then decides to participate if and only if

$$f\left(\sigma^2 \sum_{i \in S} w_i^2 + \frac{2}{\eta^2}\right) - c_i w_i \eta \geq o.$$

The platform then aims to solve the following optimization problem:

$$
\begin{aligned}
\min_{\eta, \mathbf{w}, S} \quad & \sum_{i \in S} w_i^2 \sigma^2 + \frac{2}{\eta^2} \\
\text{s.t.} \quad & c_i w_i \eta \leq f\left(\sigma^2 \sum_{i \in S} w_i^2 + \frac{2}{\eta^2}\right) - o \ \forall i \in S \\
& \sum_{i \in S} w_i = 1 \\
& w_i \geq 0 \ \forall i
\end{aligned}
\tag{1}
$$

where the first constraint ensures that agents in S choose to participate, and the last two constraints enforce that the platform's estimator is unbiased.

Finally, we make the following assumption that function f is well-behaved for our purposes:

Assumption 1. *f is concave and differentiable.*

B. The Privacy-Constrained Agent Model

We now consider a variant of our model of agent behavior. In the "privacy-constrained agent model", each agent i, on top of a privacy cost, also has a *privacy budget* B_i which is the maximum privacy cost the agent is willing to incur. An agent's utility for participation is then given by

$$
u_i(\mathbf{w}, \eta) = \begin{cases} g\left(\sum_{i \in S} w_i^2 \sigma^2 + \frac{2}{\eta^2}\right) & \text{if } c_i w_i \eta \leq B_i \\ -\infty & \text{otherwise,} \end{cases}
$$

where g is a non-negative (agents get more utility from participating than non-participating) and decreasing function. I.e., an agent is never willing to participate if his privacy budget is violated. Otherwise, if the agent's privacy requirement is met, his utility is given by a function of the accuracy of the model. The analyst's program is then given by:

$$
\begin{aligned}
\min_{\eta, \mathbf{w}, S} \quad & \sum_{i \in S} w_i^2 \sigma^2 + \frac{2}{\eta^2} \\
\text{s.t.} \quad & c_i w_i \eta \leq B_i \quad \forall i \in S \\
& \sum_{i \in S} w_i = 1 \\
& w_i \geq 0 \ \forall i
\end{aligned}
\tag{2}
$$

Note that in this case, each agent gets utility $g\left(\sum_{i \in S} w_i^2 \sigma^2 + \frac{2}{\eta^2}\right)$, and minimizing the variance of

the platform's estimate also maximizes the agents' utilities. We can further re-write the program as

$$\min_{\eta,\mathbf{w},S} \sum_{i\in S} w_i^2\sigma^2 + \frac{2}{\eta^2}$$
$$\text{s.t. } w_i\eta \le \tau_i \quad \forall i \in S \tag{3}$$
$$\sum_{i\in S} w_i = 1$$
$$w_i \ge 0 \;\forall i,$$

where $\tau_i \triangleq \frac{B_i}{c_i}$ is called the *privacy threshold* of agent i. We assume $\tau_i > 0$ for all i; agents with $\tau_i = 0$ require $w_i = 0$, do not affect the objective function, and can be dropped from the computation without loss of generality.

This model is a more tractable variant of the quasi-linear one in that participation decisions by the agents are significantly simplified. Even when the value of η used by the platform is known or announced, the "quasi-linear" model considers agents that trade-off their privacy losses with their benefit from the platform's model. In this case, an agent's participation decision depends on them being able to anticipate the quality of the final model, which requires access to the weights given to other agents. To do so, the platform either needs to communicate these weights to the agent, or each agent can solve the optimization himself, which may require unrealistic knowledge about the other agents' costs as well as unrealistic reasoning and computational power. In contrast, an agent in the "privacy-constrained" setting makes a simpler decision that only depends on his own weight w_i (this can be interpreted as a promise to the agent on how much their data is going to be used at most) and their privacy preferences τ_i. In Section VI, we will note that despite its relative simplicity, the "privacy-contrained" model offers similar insights to that of the "quasi-linear" model in Section V; this provides evidence that even this simplified model can provide valuable guidance on how to acquire and use data from agents with heterogeneous privacy preferences.

V. Characterizing the Optimal Solution under the "Quasi-Linear" Agent Model

Recall that the optimization problem solved by the platform is given by

$$\min_{\eta,\mathbf{w},S} \sum_{i\in S} w_i^2\sigma^2 + \frac{2}{\eta^2}$$
$$\text{s.t. } c_i w_i \eta \le f\left(\sigma^2\sum_{i\in S} w_i^2 + \frac{2}{\eta^2}\right) - o \;\forall i \in S \tag{4}$$
$$\sum_{i\in S} w_i = 1$$
$$w_i \ge 0 \;\forall i$$

a) Re-writing the optimization problem: We first rewrite the optimization problem solved by the analyst in a simpler form. To do so, we show how to drop the dependency of the optimization program in S. We now only need to optimize over $\vec{w}$ and η.

Claim 1. *Consider the following program:*

$$\min_{\eta,\mathbf{w}} \sigma^2\sum_{i=1}^n w_i^2 + \frac{2}{\eta^2}$$
$$\text{s.t. } c_i w_i \eta \le f\left(\sigma^2\sum_{i=1}^n w_i^2 + \frac{2}{\eta^2}\right) - o \;\forall i \in [n] \tag{5}$$
$$\sum_{i=1}^n w_i = 1$$
$$w_i \ge 0 \;\forall i.$$

Let $S = \{i \text{ s.t. } w_i > 0\}$. Then $(\mathbf{w}, \eta, S)$ is an optimal solution to Program (4) if and only if $(\mathbf{w}, \eta)$ is an optimal solution to Program (5), and both programs have the same optimal value.

Proof Sketch. Clearly $(\mathbf{w}, \eta, S)$ yields the same objective value for Program (4) as $(\mathbf{w}, \eta)$ does for Program (5). Further, $(\mathbf{w}, \eta, S)$ is feasible for Program (4) if and only if $(\mathbf{w}, \eta)$ is feasible for Program (5). Both statements put together imply that both programs have the same optimal value and that said optimal value is reached at $(\mathbf{w}, \eta, S)$ and $(\mathbf{w}, \eta)$ respectively. More details are provided in Appendix B-A. □

In short, note that if we find an optimal solution to Program 5, we can construct an optimal solution to Program (4) with the same objective value. Studying Program 5 is without loss of generality.

Note however that the above optimization problem may be hard to solve directly as it is not convex: indeed, $(\mathbf{w}, \eta) \to w_i\eta$ is not a jointly convex function of $\mathbf{w}$ and

η. To deal with this issue, we note that once we fix the value of η, the problem is now entirely convex. Indeed, i) the objective function is convex in $\mathbf{w}$, ii) $c_i w_i \eta$ is convex in $\mathbf{w}$ and $-f\left(\sigma^2 \sum_{i=1}^n w_i^2 + \frac{2}{\eta^2}\right)$ is convex in $\mathbf{w}$ (because -f is convex increasing and $\sigma^2 \sum_{i=1}^n w_i^2 + \frac{2}{\eta^2}$ is convex), and iii) the weight constraints are linear. In this section, we mostly focus on understanding this convex optimization problem for any fixed value of η. Finding the best η corresponds to finding the optimum of a one-dimensional function, which can be approximated heuristically through black-box optimization techniques.

b) Properties of the optimal solution: In the rest of this section, we order agents as a function of their privacy costs. I.e., without loss of generality, we number agents such that $c_1 \leq \ldots \leq c_n$. As mentioned above, we now consider optimization at fixed η. I.e., for any given η, we aim to solve program

$$OPT(\eta) =$$
$$\min_{\mathbf{w}} \quad \sigma^2 \sum_{i=1}^n w_i^2 + \frac{2}{\eta^2}$$
$$\text{s.t. } c_i w_i \eta \leq f\left(\sigma^2 \sum_{i=1}^n w_i^2 + \frac{2}{\eta^2}\right) - o \ \forall i \in [n]$$
$$\sum_{i=1}^n w_i = 1$$
$$w_i \geq 0 \ \forall i.$$

(6)

We first note the following simple monotonicity result on the structure of an optimal solution to Program (6) hence (5):

Claim 2. *Take any $\eta \geq 0$ such that Program (6) is feasible, then any optimal solution to Program (6) satisfies $w_1 \geq \ldots \geq w_n$.*

I.e., as we would intuitively expect, agents with smaller privacy costs get more weight in the computation. This allows the platform to provide more privacy to agents with higher costs to ensure said costs do not become too high and violate the participation constraint. We also note the following monotonicity result of independent interest, which states that the privacy costs of the agents are in fact monotone increasing in the c_i's. I.e., agents with less stringent privacy attitudes also end up incurring lower privacy costs.

Claim 3. *Any optimal solution $\mathbf{w}$ to Program (6) satisfies $c_i w_i \eta \leq c_j w_j \eta$ for all $i < j$.*

The proofs of the three previous claims are provided in Appendix B-A. The proofs are by contradiction and show that if an optimal solution satisfies the condition of each of the claim, then we can construct a feasible solution with better objective, contradicting optimality. Putting the previous claims together, we obtain the following corollary:

Corollary 1. *Any optimal solution $\mathbf{w}$ satisfies $w_i > 0$ for all $i \in [n]$.*

Proof. Suppose this were not true, i.e. for some i, $w_i = 0$. Then by monotonicity of $\mathbf{w}$ proven in Claim 2, it must be that $w_n = 0$. But then, by Claim 3, it must be that $c_i w_i \leq c_n w_n = 0$ for all i. This implies $w_i = 0$ for all i, which contradicts $\sum_i w_i = 1$. $\square$

We note that in our optimal solution, *every* agent is incentivized to participate in the platform and to report their data. This is the result of a self-reinforcing effect exhibited in our setting: on the one hand, more participation means that the platform computes a more accurate model, which incentivizes more agent participation; on the other hand, more participation lowers the privacy costs of the agents (as it lowers how much the computation depends on any given agent's data), which also helps incentivizing more participation.

c) A semi-closed form characterization: We now provide the main characterization result of this section; namely, a semi-closed form solution for Program (6).

Theorem 1. *Assume Program 6 is feasible. Let $\mathbf{w}$ be any optimal solution to Program 6. There exists constants K and W and an integer t such that $w_i = W$ for all $i \leq t$ and $w_i = K/c_i$ for all $i \geq t + 1$.*

Proof sketch. The full proof of the result relies follows by examining the implications of the Karush–Kuhn–Tucker (KKT) conditions for optimality and is provided in Appendix A-A. One technicality is that the KKT conditions require that Slater's condition holds. This means that the optimization program needs to be strictly feasible, i.e. there must exist a feasible solution such that all inequality constraints are strictly satisfied. To circumvent this issue, we note that when we do not have strict feasibility, any feasible (hence the optimal) solution must make *all* participation constraints tight hence is easy to characterize. $\square$

We remark that our optimal solution exhibits interesting structure. First, there is a pooling region in which the agents with the lowest privacy costs are given the same weights. Then, agent weights start decreasing in their

cost to ensure that their privacy losses do not become too big. We note that this result is in line with that of [2]. This is perhaps surprising given that [2] considers a different objective and constraints for the platform.

A potential explanation may be that absent privacy constraint, the optimal solution in terms of variance is to give the same weight to every agent. However, this may not be possible due to the agents' privacy requirements. Instead, one wants to have a solution that keeps the weights of different agents equal when possible to minimize the variance due to these agents, and only give a different, lower weight to the agents when this is unavoidable to ensure they participate in the computation.

d) Finding the optimal value of η: One possible approach to optimize over the value of η is to do a grid search over said 1-dimensional parameter. However, $OPT(\eta)$ is a black-box, not well understood function of η, that may be complex to optimize over. Another approach is to refine our understanding of the relationship between K, W, and η. One way to do so is to first note that if we know t, there is a closed-form relationship between K and W. In particular, we have that

$$tW + K \sum_{i>t} \frac{1}{c_i} = 1,$$

implying that

$$W = \frac{1}{t}\left(1 - K\sum_{i>t}\frac{1}{c_i}\right).$$

From the proof of Theorem 1 found in Appendix A-A, we also know that the participation constraint is tight for all agents $i > t$ with $w_i = \frac{K}{c_i}$, hence it must be that

$$K\eta = f\left(\sigma^2 t W^2 + K^2 \sum_{i>t} \frac{1}{c_i}^2 + \frac{2}{\eta^2}\right) - o.$$

This can be rewritten as

$$K\eta$$
$$= f\left(\frac{\sigma^2}{t}\left(1 - K\sum_{i>t}\frac{1}{c_i}\right)^2 + K^2\sum_{i>t}\frac{1}{c_j}^2 + \frac{2}{\eta^2}\right) - o.$$

In particular, for each possible value of t, we can restrict our search to the parameters K and η that satisfy the above equation. In the special case where f is linear, this equation is quadratic and has (at most) two well-behaved solutions that depend continuously on the value of η. This facilitates a grid search approach to find the best η for each possible value of t. We can then simply pick the value of t that leads to the best objective value.

VI. Characterizing the Optimal Solution under the "Privacy-Constrained" Agent Model

Recall that the optimization program solved by the platform is given by:

$$\begin{aligned}
\min_{\eta, \mathbf{w}, S} \quad & \sum_{i \in S} w_i^2 \sigma^2 + \frac{2}{\eta^2} \\
\text{s.t.} \quad & w_i \eta \leq \tau_i \quad \forall i \in S \\
& \sum_{i \in S} w_i = 1 \\
& w_i \geq 0 \ \forall i,
\end{aligned} \tag{7}$$

a) Re-writing the optimization problem: We start by noting that Program 7 can be rewritten in a simpler form involving no S variable. Indeed:

Claim 4. *Consider the following program:*

$$\begin{aligned}
\min_{\eta, \mathbf{w}} \quad & \sum_{i=1}^{n} w_i^2 \sigma^2 + \frac{2}{\eta^2} \\
\text{s.t.} \quad & w_i \eta \leq \tau_i \quad \forall i \in [n] \\
& \sum_{i=1}^{n} w_i = 1 \\
& w_i \geq 0 \ \forall i,
\end{aligned} \tag{8}$$

Let $S = \{i \ s.t. \ w_i > 0\}$. Then $(\mathbf{w}, \eta, S)$ is an optimal solution to Program (7) if and only if $(\mathbf{w}, \eta)$ is an optimal solution to Program (8), and both programs have the same optimal value.

Proof. The proof is nearly identical to that of Claim 1 and is omitted for the sake of brevity. □

Once again, this optimization problem is not convex. However, if we fix η and only consider $\mathbf{w}$ as a variable, our optimization problem becomes convex. We can then solve the problem efficiently for any desired value of η, then search over η to find the optimal solution.

In the rest of this section, we first show that the optimal solution has similar positivity and monotonicity properties to that of the "quasi-linear" model. We then show that we can characterize the optimal solution in semi-closed form. Finally, we exploit the structure of our problem to provide a simple characterization and algorithm for finding the optimal η.

b) Properties of the optimal solution: Without loss of generality, we number agents so that $\tau_1 \geq \tau_2 \geq \ldots \geq \tau_n$. I.e., agents with higher indices have more stringent privacy requirements. As mentioned above, we

now consider the optimization at fixed η, and study the problem

$$
\begin{aligned}
\min_{\mathbf{w}} \quad & \sum_{i=1}^{n} w_i^2 \sigma^2 + \frac{2}{\eta^2} \\
\text{s.t. } \quad & w_i \eta \leq \tau_i \quad \forall i \in [n] \\
& \sum_{i=1}^{n} w_i = 1 \\
& w_i \geq 0 \ \forall i.
\end{aligned}
\tag{9}
$$

To draw a parallel with the "quasi-linear" model, we first show that this variant of our model exhibits strong monotonicity and positivity properties.

Claim 5. *Suppose* $\mathbf{w}$ *is an optimal solution. Then* $w_i > 0 \ \forall i \in [n]$, $w_1 \geq \ldots \geq w_n$, *and* $\frac{w_1}{\tau_1} \leq \ldots \leq \frac{w_n}{\tau_n}$.

c) A semi-closed form solution: Claim 5 provides a high level understanding of the shape of the optimal solution. We now refine this understanding by providing a semi-closed form solution to Program 8.

Theorem 2. *Let* $\mathbf{w}$ *be any optimal solution to Program 9 (assuming feasibility). There exists* $t \in \{0, \ldots, n\}$ *and* $W \geq 0$ *such that is given by* $w_i = W$ *for all* $i \leq t$ *and* $w_i = \frac{\tau_i}{\eta}$ *for all* $i \geq t + 1$.

Proof Sketch. As before, the full proof of the result relies on the Karush–Kuhn–Tucker (KKT) conditions and is provided in Appendix A-B. The proof suffers from the same technicality that the KKT conditions require that the optimization program is strictly feasible, and we use the same techniques as for Theorem 1 to circumvent said issue. □

We note that the above result bears similarities to that of Theorem 1 in the "quasi-linear" model. Indeed, note that τ_i is a parameter that is smaller as the privacy preferences of agent i are more stringent, similarly to $1/c_i$ in the "quasi-linear" model. Both solutions then have the same structure: agents with more lax privacy requirements are pooled together and have the same weight, while agents with stronger requirements see their weight decrease as a function of how strong that requirement is.

Corollary 2. *There exists* $t \in \{0, \ldots, n\}$ *such that the optimal solution to Program 9 (assuming feasibility) is given by* $w_i = \frac{1}{t}\left(1 - \frac{1}{\eta} \sum_{i=t+1}^{n} \tau_i\right)$ *for all* $i \leq t$ *and* $w_i = \frac{\tau_i}{\eta}$ *for all* $i \geq t + 1$.

Proof. It suffices to use the fact that the weights of the agents must sum to 1. I.e., $\sum_{i=1}^{n} w_i = 1$ can be rewritten as

$$
\sum_{i=1}^{t} W + \sum_{i=t+1}^{n} \frac{\tau_i}{\eta} = 1,
$$

or equivalently

$$
tW + \frac{1}{\eta} \sum_{i=t+1}^{n} \tau_i = 1.
$$

This immediately leads to $W = \frac{1}{t}\left(1 - \frac{1}{\eta} \sum_{i=t+1}^{n} \tau_i\right)$. □

d) Finding the optimal value of η exactly: We show that, in fact, η can be found by simply minimizing a function of a single variable on a closed interval:

Claim 6. *The optimal value of η is given by*

$$
\eta^* = \arg\min_{\eta} h(\eta) \ \text{s.t. } \eta \in \left[\sum_{i=1}^{t+1} \tau_i, t\tau_t + \sum_{i=1}^{t+1} \tau_i\right],
\tag{10}
$$

where

$$
h(\eta) = \frac{\sigma^2}{t^2} \sum_{i=1}^{t} \left(1 - \frac{1}{\eta} \sum_{i=t+1}^{n} \tau_i\right)^2 + \frac{\sigma^2}{\eta^2} \sum_{i=t+1}^{n} \tau_i^2 + \frac{2}{\eta^2}.
$$

Proof. Assuming the optimal value of t is known, plugging the solution of Corollary 2 back into Program (7) shows that η must solve

$$
\begin{aligned}
\min_{\eta \geq 0} \quad & \frac{\sigma^2}{t^2} \sum_{i=1}^{t} \left(1 - \frac{1}{\eta} \sum_{i=t+1}^{n} \tau_i\right)^2 + \frac{\sigma^2}{\eta^2} \sum_{i=t+1}^{n} \tau_i^2 + \frac{2}{\eta^2} \\
\text{s.t. } \quad & \frac{1}{t}\left(1 - \frac{1}{\eta} \sum_{i=t+1}^{n} \tau_i\right) \eta \leq \tau_i \quad \forall i \leq t \\
& \frac{1}{t}\left(1 - \frac{1}{\eta} \sum_{i=t+1}^{n} \tau_i\right) \geq 0
\end{aligned}
$$

Note that we dropped the constraint that the weights sum to 1: this is guaranteed to hold for any plugged-in solution of the form given in Corollary 2. Using the fact that $\tau_1 \geq \ldots \geq \tau_t$, we can rewrite the problem as

$$
\begin{aligned}
\min_{\eta \geq 0} \quad & \frac{\sigma^2}{t} \left(1 - \frac{1}{\eta} \sum_{i=t+1}^{n} \tau_i\right)^2 + \frac{\sigma^2}{\eta^2} \sum_{i=t+1}^{n} \tau_i^2 + \frac{2}{\eta^2} \\
\text{s.t. } \quad & \frac{1}{t}\left(1 - \frac{1}{\eta} \sum_{i=t+1}^{n} \tau_i\right) \eta \leq \tau_t, \\
& \frac{1}{t}\left(1 - \frac{1}{\eta} \sum_{i=t+1}^{n} \tau_i\right) \geq 0,
\end{aligned}
\tag{11}
$$

9

or equivalently

$$\min_{\eta} \ \frac{\sigma^2}{t}\left(1 - \frac{1}{\eta}\sum_{i=t+1}^{n}\tau_i\right)^2 + \frac{\sigma^2}{\eta^2}\sum_{i=t+1}^{n}\tau_i^2 + \frac{2}{\eta^2}$$

$$\text{s.t. } \eta \le t\tau_t + \sum_{i=t+1}^{n}\tau_i, \qquad (12)$$

$$\eta \ge \sum_{i=1}^{t+1}\tau_i.$$

$\square$

We further show that this is a simple optimization problem in that $f(\eta)$ is well-behaved and easy to minimize.

Claim 7. *There exists η^* such that $f(\eta)$ is decreasing for $\eta < \eta^*$, and increasing for $\eta > \eta^*$. In turn, η^* is the unique solution to $f'(\eta) = 0$, and is given in closed form by*

$$\eta^* = \frac{\left(\sum_{i=t+1}^{n}\tau_i\right)^2 + t\sum_{i=t+1}^{n}\tau_i^2 + \frac{2t}{\sigma^2}}{\sum_{i=t+1}^{n}\tau_i}.$$

If $\eta^* \in [\sum_{i=1}^{t+1}\tau_i, t\tau_t + \sum_{i=1}^{t+1}\tau_i]$, it minimizes f; otherwise, the minimizer is either $\sum_{i=1}^{t+1}\tau_i$ or $t\tau_t + \sum_{i=1}^{t+1}\tau_i$.

The proof follows from simple algebra and is given in Appendix B-B2. We note that the above characterization gives us an immediate algorithm to find the optimal η. Indeed, it suffices to explore all n possible values of t. For each η, then, one only has to compare $f(\eta^*)$, $f\left(\sum_{i=1}^{t+1}\tau_i\right)$ and $f\left(t\tau_t + \sum_{i=1}^{t+1}\tau_i\right)$. This algorithm takes time $O(n)$.

e) Incentive properties of the "privacy-constrained" model: Finally, we remark that the privacy-constrained model enjoys nice incentive properties: e.g., it is a weakly dominated strategy for agents to misreport their privacy thresholds. We note that this property holds *without having to pay agents to report their privacy preferences truthfully*. This is a major advantage in that it reflects what happens in real-life platforms, who often do not pay their users to incentivize participation; in fact, platforms commonly ask users to pay to be able to access the service they offer in return. We divide the incentive properties in the following two claims:

Claim 8. *For any agent i, reporting $\hat{\tau}_i < \tau_i$ is a weakly dominated strategy.*

Proof. Fix the participation strategy of all the other agents–let S_{-i} the set of agents that decide to join the

platform and report their data–, as well as their reports $\hat{\tau}_j$ for all $j \in S_{-i}$. Suppose the set of participating agents is $S = S_{-i} \cup i$ (i.e. i decides to participate) and $\hat{\tau}_i < \tau_i$. Let $OPT(S,\hat{\tau})$ be the optimal objective value of Program (8) when the inputs are $S, \hat{\tau}$. We immediately have that

$$OPT(S,\hat{\tau}) \ge OPT(S,(\tau_i, \hat{\tau}_{-i}))$$

by virtue of the left-hand side optimization program being strictly more constrained. Since agent i's privacy constraint is always satisfied when reporting his true threshold (by construction of Program (8)), agent i gets utility $g(OPT(S,(\tau_i, \hat{\tau}_{-i})))$. This is at least as high (by virtue of g being decreasing) as the utility agent i gets from misreporting as above, since then agent i gets utility either $g(OPT(S,\hat{\tau}))$ or $-\infty$ if his true privacy constraint is not satisfied. $\square$

Claim 9. *Fix any agent i. Reporting $\hat{\tau}_i > \tau_i$ cannot increase agent i's utility.*

Proof. Let us once again fix S_{-i} and τ_{-i}. We have two cases for agent i:

- In $(S,\hat{\tau})$, agent i receives $w_i\eta \le \tau_i$. In this case, note that the optimization program with threshold τ_i and $\hat{\tau}_i$ are equivalent and

$$OPT(S,\hat{\tau}) = OPT(S,(\tau_i, \hat{\tau}_{-i})).$$

  In this case, agent i's privacy constraint is satisfied and he gets utility $g(OPT(S,\hat{\tau})) = g(OPT(S,(\tau_i, \hat{\tau}_{-i})))$. I.e., his utility is unchanged.
- Otherwise, agent i's privacy constraint is not satisfied and he receives the worst possible utility of $-\infty$.

This concludes the proof. $\square$

REFERENCES

[1] A. Fallah, A. Makhdoumi, A. Malekian, and A. Ozdaglar, "Optimal and differentially private data acquisition: Central and local mechanisms," in *Proceedings of the 2022 ACM Conference on Economics and Computation*, 2022, p. 1141.

[2] Y. Chen, N. Immorlica, B. Lucier, V. Syrgkanis, and J. Ziani, "Optimal data acquisition for statistical estimation," in *Proceedings of the 2018 ACM Conference on Economics and Computation*, 2018, pp. 27–44.

[3] D. Bergemann and A. Bonatti, "Markets for information: An introduction," *Annual Review of Economics*, vol. 11, pp. 85–107, 2019.

[4] A. Agarwal, M. Dahleh, and T. Sarkar, "A marketplace for data: An algorithmic solution," in *Proceedings of the 2019 ACM Conference on Economics and Computation*, 2019, pp. 701–726.

[5] A. Roth and G. Schoenebeck, "Conducting truthful surveys, cheaply," in *Proceedings of the 13th ACM Conference on Electronic Commerce*, 2012, pp. 826–843.

[6] Y. Chen and S. Zheng, "Prior-free data acquisition for accurate statistical estimation," in *Proceedings of the 2019 ACM Conference on Economics and Computation*, 2019, pp. 659–677.

[7] D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar, "Too much data: Prices and inefficiencies in data markets," National Bureau of Economic Research, Tech. Rep., 2019.

[8] G. Liao, Y. Su, J. Ziani, A. Wierman, and J. Huang, "The privacy paradox and optimal bias-variance trade-offs in data acquisition," *ACM SIGMETRICS Performance Evaluation Review*, vol. 49, no. 2, pp. 6–8, 2022.

[9] V. Gkatzelis, C. Aperjis, and B. A. Huberman, "Pricing private data," *Electronic Markets*, vol. 25, no. 2, pp. 109–123, 2015.

[10] J. Abernethy, Y. Chen, C.-J. Ho, and B. Waggoner, "Low-cost learning via active data procurement," in *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, 2015, pp. 619–636.

[11] Y. Cai, C. Daskalakis, and C. Papadimitriou, "Optimum statistical estimation with strategic data sources," in *Conference on Learning Theory*. PMLR, 2015, pp. 280–296.

[12] Y. Liu and Y. Chen, "Learning to incentivize: Eliciting effort via output agreement," *arXiv preprint arXiv:1604.04928*, 2016.

[13] ——, "Sequential peer prediction: Learning to elicit effort using posted prices," in *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.

[14] Y. Liu, J. Wang, and Y. Chen, "Surrogate scoring rules," in *Proceedings of the 21st ACM Conference on Economics and Computation*, 2020, pp. 853–871.

[15] Y. Chen, Y. Liu, and C. Podimata, "Learning strategy-aware linear classifiers," *Advances in Neural Information Processing Systems*, vol. 33, pp. 15265–15276, 2020.

[16] J. Perote and J. Perote-Pena, "The impossibility of strategy-proof clustering," *Economics Bulletin*, vol. 4, no. 23, pp. 1–9, 2003.

[17] ——, "The impossibility of strategy-proof clustering," *Economics Bulletin*, vol. 4, no. 23, pp. 1–9, 2003.

[18] O. Dekel, F. Fischer, and A. D. Procaccia, "Incentive compatible regression learning," *Journal of Computer and System Sciences*, vol. 76, no. 8, pp. 759–777, 2010.

[19] R. Meir and J. S. Rosenschein, "Strategyproof classification," *ACM SIGecom Exchanges*, vol. 10, no. 3, pp. 21–25, 2011.

[20] R. Meir, A. D. Procaccia, and J. S. Rosenschein, "Algorithms for strategyproof classification," *Artificial Intelligence*, vol. 186, pp. 123–156, 2012.

[21] L. K. Fleischer and Y.-H. Lyu, "Approximately optimal auctions for selling privacy when costs are correlated with data," in *Proceedings of the 13th ACM conference on electronic commerce*, 2012, pp. 568–585.

[22] K. Nissim, C. Orlandi, and R. Smorodinsky, "Privacy-aware mechanism design," in *Proceedings of the 13th ACM Conference on Electronic Commerce*, 2012, pp. 774–789.

[23] R. Cummings, K. Ligett, A. Roth, Z. S. Wu, and J. Ziani, "Accuracy for sale: Aggregating data with a variance constraint," in *Proceedings of the 2015 conference on innovations in theoretical computer science*, 2015, pp. 317–324.

[24] R. Cummings, V. Feldman, A. McMillan, and K. Talwar, "Mean estimation with user-level privacy under data heterogeneity," in *NeurIPS 2021 Workshop Privacy in Machine Learning*, 2021.

[25] A. Ghosh and A. Roth, "Selling privacy at auction," *Games and Economic Behavior*, vol. 91, pp. 334–346, 2015.

[26] A. Ghosh, K. Ligett, A. Roth, and G. Schoenebeck, "Buying private data without verification," in *Proceedings of the fifteenth ACM conference on Economics and computation*, 2014, pp. 931–948.

[27] R. Cummings, S. Ioannidis, and K. Ligett, "Truthful linear regression," in *Conference on Learning Theory*. PMLR, 2015, pp. 448–483.

[28] G. Liao, X. Chen, and J. Huang, "Social-aware privacy-preserving mechanism for correlated data," *IEEE/ACM Transactions on Networking*, vol. 28, no. 4, pp. 1671–1683, 2020.

[29] A. Ghosh and K. Ligett, "Privacy and coordination: Computing on databases with endogenous participation," in *Proceedings of the fourteenth ACM conference on Electronic commerce*, 2013, pp. 543–560.

[30] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of cryptography conference*. Springer, 2006, pp. 265–284.

[31] C. Dwork, A. Roth *et al.*, "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.

APPENDIX A
PROOFS OF THE MAIN THEOREMS

A. Proof of Theorem 1

First, we consider the case in which Program 5 is not strictly feasible; i.e., there exists no feasible weights $\mathbf{w}$ such that for all i,

$$c_i w_i \eta < f\left(\sigma^2 \sum_{i=1}^n w_i^2 + \frac{2}{\eta^2}\right) - o.$$

Claim 10. *Suppose Program 5 is feasible, but has no strictly feasible solution. Then any feasible, hence the optimal solution is given by $w_i = K/c_i$ for all i and for some constant K.*

Proof. First, suppose the program is not strictly feasible. In particular, let us look at $\mathbf{w}$ the optimal solution to the program. We consider two cases:

1) There exists i such that

$$c_i w_i \eta < f\left(\sigma^2 \sum_{i=1}^n w_i^2 + \frac{2}{\eta^2}\right) - o.$$

Let $\tilde{w}_i = w_i + (n-1)\varepsilon$ and $\tilde{w}_j = w_j - \varepsilon$ for all $j \neq i$. Note that for ε small enough, the $\tilde{\mathbf{w}}$ define proper weights: they are positive, since by Corollary 1, $w_j > 0$ for all j, and they sum to 1 by construction. Further,

$$c_j \tilde{w}_j \eta < c_j w_j \eta \leq f\left(\sigma^2 \sum_{i=1}^n w_i^2 + \frac{2}{\eta^2}\right) - o \ \forall j \neq i,$$

hence the participation constraints are strict for all $j \neq i$. Finally, for ε small enough, by continuity of f, we have

$$c_i \tilde{w}_i \eta < f\left(\sigma^2 \sum_{i=1}^n \tilde{w}_i^2 + \frac{2}{\eta^2}\right) - o,$$

since the left-hand side converges to $c_i w_i \eta$ and the right-hand side to $f\left(\sigma^2 \sum_{i=1}^n w_i^2 + \frac{2}{\eta^2}\right) - o$ when

ε goes to 0. Hence, $\tilde{\mathbf{w}}$ is a strictly feasible solution, which is a contradiction.

2) For all i, the participation constraint is tight, i.e.

$$c_i w_i \eta = f\left(\sigma^2 \sum_{i=1}^n w_i^2 + \frac{2}{\eta^2}\right) - o.$$

Then $w_i = K/c_i$ for all i, where $K \triangleq \frac{f\left(\sigma^2 \sum_{i=1}^n w_i^2 + \frac{2}{\eta^2}\right) - o}{\eta}$ is the same for all agents i. $\qquad \square$

Now, we can consider without loss of generality the case in which the program is strictly feasible. In this case, Slater's condition holds and we can apply the KKT conditions. Note that the Lagrangian of Program (6) is given by

$$\mathcal{L}(\mathbf{w}, \vec{\lambda}, \vec{\lambda^0}, \gamma)$$
$$= \sigma^2 \sum_{i=1}^n w_i^2 + \frac{2}{\eta^2} + \gamma\left(1 - \sum_{i=1}^n w_i\right) - \sum_{i=1}^n \lambda_i^0 w_i$$
$$+ \sum_{i=1}^n \lambda_i \left(c_i w_i \eta - f\left(\sigma^2 \sum_{j=1}^n w_j^2 + \frac{2}{\eta^2}\right) + o\right)$$

The first order condition, taking the derivative with respect to w_i (remembering that f is differentiable by assumption), is given by

$$0 = 2\sigma^2 w_i + \lambda_i c_i \eta - \gamma - \lambda_i^0$$
$$- 2\sigma^2 \left(\sum_j \lambda_j\right) w_i \cdot f'\left(\sigma^2 \sum_{j=1}^n w_j^2 + \frac{2}{\eta^2}\right)$$

This implies $w_i = \frac{-\lambda_i c_i \eta + \gamma + \lambda_i^0}{2\sigma^2\left(1 - \left(\sum_j \lambda_j\right) \cdot f'\left(\sigma^2 \sum_{j=1}^n w_j^2 + \frac{2}{\eta^2}\right)\right)}$.

Since $w_i > 0$ for all i, complementary slackness yields that $\lambda_i^0 = 0$ for all i. The first order condition then simplifies to

$$w_i = \frac{-\lambda_i c_i \eta + \gamma}{2\sigma^2\left(1 - \left(\sum_j \lambda_j\right) \cdot f'\left(\sigma^2 \sum_{j=1}^n w_j^2 + \frac{2}{\eta^2}\right)\right)}.$$

We now have two cases:

1) Either $c_i w_i \eta = f\left(\sigma^2 \sum_{i=1}^n w_i^2 + \frac{2}{\eta^2}\right) - o$, i.e. the participation constraint is tight. Then we can write

$$w_i = \frac{f\left(\sigma^2 \sum_{i=1}^n w_i^2 + \frac{2}{\eta^2}\right) - o}{\eta c_i} \triangleq \frac{K}{c_i},$$

where K is a constant in that it is the same for all agents.

2) Either w_i is such that i's participation constraint is not tight. Then, by complementary slackness, we have that $\lambda_i = 0$, hence we can rewrite

$$w_i$$
$$= \frac{\gamma}{2\sigma^2\left(1 - \left(\sum_j \lambda_j\right) \cdot f'\left(\sigma^2 \sum_{j=1}^n w_j^2 + \frac{2}{\eta^2}\right)\right)}$$
$$\triangleq W$$

is a constant that is the same for all agents.

Therefore, in any optimal solution, there exists constants K and W such that either $w_i = K/c_i$, or $w_i = W$. To conclude the proof, suppose that $i < j$, but w_i satisfies case (1) above (and $w_i = K/c_i$) while w_j satisfies case (2) (and $w_j = W$). We have that $w_j < K/c_j$ since the participation constraint $c_j w_j \leq K$ is not tight for j by definition of case (2). Hence, $c_i w_i = K$ while $c_j w_j < K$, implying that $c_i w_i > c_j w_j$. This contradicts Claim 3. Therefore, for any j, if $w_j = W$, it must be that $w_i = W$ for all $i < j$. This concludes the proof.

B. Proof of Theorem 2

First, we consider the case in which Slater's condition does not hold, and there exists at least one i such that $w_i \eta = \tau_i$ in any feasible solution. We have two cases:

1) There exists j such that $w_j \eta < \tau_j$. Then let $\tilde{\mathbf{w}}$ be such that $\tilde{w}_i = w_i - \varepsilon$ for all $i \neq j$ and let $\tilde{w}_j = w_j + (n-1)\varepsilon$. When ε is small enough, $\tilde{\mathbf{w}} \geq 0$ (noting that we have $\mathbf{w} > 0$ at an optimal solution by Claim 5), the weights sum to 1, $\tilde{w}_i \eta < (w_i \eta \leq)\tau_i$ for all $i \neq j$, and $\tilde{w}_j \eta < \tau_j$. Hence $\tilde{w}$ is strictly feasible, which is a contradiction.

2) For all i, $w_i \eta = \tau_i$. Then the optimal solution is fully determined by these equations, and satisfies $w_i = \frac{\tau_i}{\eta}$ for all i.

Now, suppose we have strict feasibility, i.e. Slater's condition holds. The Lagrangian of the optimization problem is given by

$$\mathcal{L}(\mathbf{w}, \lambda, \lambda^0, \mu) = \sum_i w_i^2 \sigma^2 + \frac{2}{\eta^2} + \sum_i \lambda_i (w_i \eta - \tau_i)$$
$$+ \mu\left(\sum_i w_i - 1\right) - \sum_i \lambda_i^0 w_i.$$

The first order condition (with respect to agent i) is then given by

$$2w_i \sigma^2 + \lambda_i \eta + \mu - \lambda_i^0 = 0,$$

which implies

$$w_i = \frac{\lambda_i \eta + \mu - \lambda_i^0}{2\sigma^2}.$$

By Claim 5, $w_i > 0$, hence by KKT conditions, $\lambda_i^0 = 0$. Therefore, we can rewrite

$$w_i = \frac{\lambda_i \eta + \mu}{2\sigma^2}.$$

We now have two cases:

1) Either agent i's privacy constraint is tight. Then, $w_i \eta = \tau_i$, i.e. $w_i = \frac{\tau_i}{\eta}$.
2) Otherwise, the privacy constraint is not tight. Then, by the KKT conditions, it must be that $\lambda_i = 0$, hence $w_i = \frac{\mu}{2\sigma^2} = W$ for some constant W. Since the privacy constraint is not tight, we have in particular that $W < \frac{\tau_i}{\eta}$.

To complete the proof, suppose that we have $i < j$ such that i is in case (1) and $w_i = \tau_i/\eta$, while j is in case (2) and $w_j = W$. We have that $w_i/\tau_i \le w_j/\tau_j$ by Claim 5. This then implies that $1/\eta \le W/\tau_j$, and in turn that $\tau_j \le W\eta < \tau_j$ (remember that since j is in case (2), $W\eta < \tau_j$). This is a contradiction. Hence, it must be the case that if if w_i is in case (1), we must have $w_j = w_j/\eta$ for all subsequent $j > i$.

APPENDIX B
PROOFS OF SUPPORTING CLAIMS

A. In the Quasi-Linear Utility Model

1) Proof of Claim 1: Pick any $\mathbf{w}, \eta$, and let $S = \{i \text{ s.t. } w_i > 0\}$. We first note that

$$\sigma^2 \sum_{i \in S} w_i^2 + \frac{2}{\eta^2} = \sigma^2 \sum_{i \in [n]} w_i^2 + \frac{2}{\eta^2}$$

since $w_i = 0$ for all $i \notin S$. Hence, $(\mathbf{w}, \eta, S)$ achieves the same objective value for Program 4 as $(\mathbf{w}, \eta)$ for Program 5 for any $\mathbf{w}, \eta$, and S constructed as above.

Second, with respect to Program 4, we have that:

1) $\sum_{i \in S} w_i = 1 \Leftrightarrow \sum_{i \in [n]} w_i = 1$ by virtue of $w_i = 0$ for all $i \notin S$.
2) Since we have that

$$f\left(\sigma^2 \sum_{i \in [n]} w_i^2 + \frac{2}{\eta^2}\right) = f\left(\sigma^2 \sum_{i \in S} w_i^2 + \frac{2}{\eta^2}\right),$$

for all i, $c_i w_i \eta \le f\left(\sigma^2 \sum_{i \in [n]} w_i^2 + \frac{2}{\eta^2}\right) - o$ if and only if $c_i w_i \eta \le f\left(\sigma^2 \sum_{i \in S} w_i^2 + \frac{2}{\eta^2}\right) - o$. Therefore $(\mathbf{w}, \eta, S)$ is feasible for Program 4 if and only $(\mathbf{w}, \eta)$ is feasible for Program 5.

This is enough to conclude the proof. Indeed, since $(\mathbf{w}, \eta)$ feasible for Program (5)$(\mathbf{w}, \eta, S)$ implies $(\mathbf{w}, \eta, S)$ feasible for Program 4 and they both have the same objective value, the optimal value of Program (4)

is at least that of Program (5). Vice-versa, the optimal value of Program 5 is at least that of Program 4. Hence, Program 4 and Program 5 have the same optimal value. Further, if $(\mathbf{w}, \eta)$ is optimal, then $(\mathbf{w}, \eta, S)$ is optimal by virtue of having the same objective value, and vice-versa. This concludes the proof.

2) Proof of Claim 2: Let $\mathbf{w}$ be an optimal solution to Program (6). Suppose there exists $i < j$ such that $w_i < w_j$. Now, let us look at a possible alternative solution $\tilde{\mathbf{w}}$ where $\tilde{w}_i \triangleq w_i + \varepsilon$, $\tilde{w}_j \triangleq w_j - \varepsilon$ for $\varepsilon > 0$ small enough, and $\tilde{w}_k \triangleq w_k$ for all agents $k \ne i, j$. We will show that this solution leads to a smaller objective, contradicting optimality.

First,

$$\sum_{k=1}^{n} \tilde{w}_k^2 - \sum_{k=1}^{n} w_k^2 = (w_i + \varepsilon)^2 - w_i^2 + (w_j - \varepsilon)^2 - w_j^2$$

$$= 2w_i \varepsilon + \varepsilon^2 + (\varepsilon^2 - 2w_j \varepsilon)$$

$$= 2\varepsilon(w_i - w_j + \varepsilon)$$

$$< 0,$$

for ε small enough, as $w_i < w_j$. This shows that $\tilde{\mathbf{w}}$ leads to a better variance than $\mathbf{w}$, since it directly implies

$$\sum_{k=1}^{n} \tilde{w}_k^2 \sigma^2 + \frac{2}{\eta^2} < \sigma^2 \sum_{k=1}^{n} w_i^2 + \frac{2}{\eta^2}.$$

Second, since $c_i \le c_j$, $\tilde{w}_j < w_j$, for ε small enough we have $\tilde{w}_i \le \tilde{w}_j$, and because f is decreasing, it follows that

$$c_i \tilde{w}_i \eta \le c_i \tilde{w}_j \eta < c_j w_j \eta \le f\left(\sigma^2 \sum_{i=1}^{n} w_i^2 + \frac{2}{\eta^2}\right)$$

$$\le f\left(\sigma^2 \sum_{i=1}^{n} \tilde{w}_i^2 + \frac{2}{\eta^2}\right).$$

Further, because $\tilde{w}_j < w_j$, we have

$$c_j \tilde{w}_j \eta < c_j w_j \eta \le f\left(\sigma^2 \sum_{i=1}^{n} w_i^2 + \frac{2}{\eta^2}\right)$$

$$\le f\left(\sigma^2 \sum_{i=1}^{n} \tilde{w}_i^2 + \frac{2}{\eta^2}\right).$$

Therefore, $\tilde{\mathbf{w}}$ is feasible, since it satisfies the participation constraints and that the weights are still positive and sum to 1. This concludes the proof.

3) Proof of Claim 3: Let $\mathbf{w}$ be an optimal solution with $c_i w_i \eta > c_j w_j \eta$; in particular, it must be that $w_i > w_j$ since $c_i \le c_j$. For small enough ε, let $\tilde{w}_i = w_i - \varepsilon$, $\tilde{w}_j = w_j + \varepsilon$, and $\tilde{w}_k = w_k$ i for any other agent $k \ne$

i, j. First, we note that this transformation decreases the variance. Indeed,

$$
\begin{aligned}
& w_i^2 + w_j^2 - \tilde{w}_i^2 - \tilde{w}_j^2 \\
&= w_i^2 + w_j^2 - w_i^2 + 2\varepsilon w_i - \varepsilon^2 - w_j^2 - 2\varepsilon w_j - \varepsilon^2 \\
&= 2\varepsilon(w_i - w_j) - 2\varepsilon^2 \\
&= 2\varepsilon(w_i - w_j - \varepsilon) \\
&> 0
\end{aligned}
$$

when ε is small enough, by virtue of $w_i > w_j$. Further, the constraints that the weights must sum to 1 still holds, as well as the non-negativity constraint so long as ε is small enough (smaller than w_i). Finally, $c_i \tilde{w}_i \eta = c_i(w_i - \varepsilon)\eta \le c_i w_i \eta$, and $c_j \tilde{w}_j \eta = c_j(w_j \eta + \varepsilon \eta) \le c_i w_i \eta$ for small enough ε (as $c_j w_j \eta < c_i w_i \eta$); combining this with the fact that the variance decreases, the participation constraints still holds. Therefore, $\tilde{\mathbf{w}}$ is feasible for Program 5 and has strictly better objective value than an optimal solution, which is a contradiction.

B. In the Privacy-Constrained Utility Model

1) Proof of Claim 5: We show the results in the claim by contradiction. First, suppose there exists i such that $w_i = 0$. We will show that we can construct alternative weight vector $\mathbf{w}$ that is feasible and leads to a strictly better objective value, contradicting optimality of $\mathbf{w}$. To do so, let j be such that $w_j > 0$, and let $\tilde{w}_i = \varepsilon$, $\tilde{w}_j = w_j - \varepsilon$, and $\tilde{w}_j = w_j$ for all $k \neq i, j$. For ε small enough, note that $\tilde{\mathbf{w}}$ is feasible: all weights remain non-negative, sum to 1, $\tilde{w}_j \eta \le w_j \eta \le \tau_j$, and $\tilde{w}_i \eta = \varepsilon \eta \le \tau_i$ so as long as ε is sufficiently small. Further, the objective value under $\tilde{\mathbf{w}}$ is smaller than under $\mathbf{w}$. Indeed, the change in variance (renormalized by $1/\sigma^2$) is given by

$$
\begin{aligned}
w_i^2 - \tilde{w}_i^2 + w_j^2 - \tilde{w}_j^2 &= -\varepsilon^2 + w_j^2 - (w_j - \varepsilon)^2 \\
&= -\varepsilon^2 + w_j^2 - w_j^2 + 2w_j \varepsilon - \varepsilon^2 \\
&= 2\varepsilon(w_j - \varepsilon) \\
&> 0.
\end{aligned}
$$

where the last inequality follows from ε being small enough. This is a contradiction.

Second, suppose there exists $i < j$ such that $w_i < w_j$. Let us construct as before an alternative weight vector $\tilde{\mathbf{w}}$ such that $\tilde{w}_i = w_i + \varepsilon$, $\tilde{w}_j = w_j - \varepsilon$, and $\tilde{w}_k = w_k$ for all $k \neq i, j$. First, $\tilde{\mathbf{w}}$ is feasible: $\sum_i \tilde{w}_i = 1$, $\tilde{w}_j \ge 0$ for ε small enough (since $w_j > 0$), $\tilde{w}_i \eta < w_j \eta \le \tau_j \le \tau_i$

for ε small enough (as $w_i < w_j$), and $\tilde{w}_j \eta \le w_j \eta \le \tau_j$. Further, $\tilde{\mathbf{w}}$ yields better variance than $\mathbf{w}$. Indeed,

$$
\begin{aligned}
& \sum_{i=1}^{n} w_i^2 - \sum_{i=1}^{n} \tilde{w}_i^2 \\
&= w_i^2 - \tilde{w}_i^2 + w_j^2 - \tilde{w}_j^2 \\
&= w_i^2 - (w_i + \varepsilon)^2 + w_j^2 - (w_j - \varepsilon)^2 \\
&= w_i^2 - w_i^2 - \varepsilon^2 - 2w_i\varepsilon + w_j^2 - w_j^2 + 2w_j\varepsilon - \varepsilon^2 \\
&= 2(w_j - w_i - \varepsilon) \\
&> 0
\end{aligned}
$$

for ε small enough, remembering that $w_j > w_i$. This is a contradiction.

Finally, suppose there exists $i < j$ such that $\frac{w_i}{\tau_i} > \frac{w_j}{\tau_j}$ (note that since $\tau_i \ge \tau_j$, this also implies $w_i > w_j$). Then, consider alternative weight vector $\tilde{w}_i = w_i - \varepsilon$, $\tilde{w}_j = w_j + \varepsilon$, and $\tilde{w}_k = w_k$ for all $k \neq i, j$. First, we note that $\mathbf{w}$ is feasible. Indeed, $\sum_i \tilde{w}_i = 1$, $\tilde{w}_i \eta < w_i \eta \le \tau_i$, and for ε small enough,

$$
\frac{\tilde{w}_j}{\tau_j} \eta < \frac{\tilde{w}_i}{\tau_i} \eta < \frac{w_i}{\tau_i} \eta \le 1.
$$

$\tilde{\mathbf{w}}$ also has lower variance than $\mathbf{w}$, by a similar calculation as before:

$$
\begin{aligned}
& \sum_{i=1}^{n} w_i^2 - \sum_{i=1}^{n} \tilde{w}_i^2 \\
&= w_i^2 - \tilde{w}_i^2 + w_j^2 - \tilde{w}_j^2 \\
&= w_i^2 - (w_i - \varepsilon)^2 + w_j^2 - (w_j + \varepsilon)^2 \\
&= w_i^2 - w_i^2 - \varepsilon^2 + 2w_i\varepsilon + w_j^2 - w_j^2 - 2w_j\varepsilon - \varepsilon^2 \\
&= 2\varepsilon(w_i - w_j - \varepsilon) \\
&> 0
\end{aligned}
$$

where the last step follows from $w_i > w_j$ and ε small enough. This is a contradiction.

2) Proof of Claim 7: We have that

$$
\begin{aligned}
& f'(\eta) \\
&= \frac{2\sigma^2 \sum_{i=t+1}^{n} \tau_i}{t\eta^2}\left(1 - \frac{1}{\eta}\sum_{i=t+1}^{n}\tau_i\right) - \frac{2\sigma^2}{\eta^3}\sum_{i=t+1}^{n}\tau_i^2 - \frac{4}{\eta^3} \\
&= \frac{2\sigma^2 \sum_{i=t+1}^{n} \tau_i}{t\eta^3}\left(\eta - \sum_{i=t+1}^{n}\tau_i\right) - \frac{2\sigma^2}{\eta^3}\sum_{i=t+1}^{n}\tau_i^2 - \frac{4}{\eta^3} \\
&= \frac{2\sigma^2}{\eta^3}\left(\frac{\eta}{t}\sum_{i=t+1}^{n}\tau_i - \frac{1}{t}\left(\sum_{i=t+1}^{n}\tau_i\right)^2 - \sum_{i=t+1}^{n}\tau_i^2 - \frac{2}{\sigma^2}\right).
\end{aligned}
$$

In turn, $f'(\eta) < 0$ if and only if $\eta < \frac{\left(\sum_{i=t+1}^{n}\tau_i\right)^2 + t\sum_{i=t+1}^{n}\tau_i^2 + 2t/\sigma^2}{\sum_{i=t+1}^{n}\tau_i}$ and $f'(\eta) > 0$ if and

only if $\eta < \frac{\left(\sum_{i=t+1}^{n} \tau_i\right)^2 + t\sum_{i=t+1}^{n} \tau_i^2 + 2/\sigma^2}{\sum_{i=t+1}^{n} \tau_i}$. Finally, note that $f'(\eta) = 0$ can be written as

$$\frac{\eta}{t} \sum_{i=t+1}^{n} \tau_i - \frac{1}{t} \left(\sum_{i=t+1}^{n} \tau_i \right)^2 - \sum_{i=t+1}^{n} \tau_i^2 - \frac{2}{\sigma^2} = 0.$$

This immediately leads to

$$\eta^* = \frac{\left(\sum_{i=t+1}^{n} \tau_i \right)^2 + t \sum_{i=t+1}^{n} \tau_i^2 + \frac{2t}{\sigma^2}}{\sum_{i=t+1}^{n} \tau_i}.$$