X

Privacy and coordination: Computing on databases with endogenous participation

Arpita Ghosh, Cornell University
Katrina Ligett, Caltech

We propose a simple model where individuals in a privacy-sensitive population decide whether or not to participate in a pre-announced noisy computation by an analyst, so that the database itself is *endogenously* determined by individuals' participation choices. The privacy an agent receives depends both on the announced noise level, *as well as* how many agents choose to participate in the database. Each agent has some minimum privacy requirement, and decides whether or not to participate based on how her privacy requirement compares against her expectation of the privacy she will receive if she participates in the computation. This gives rise to a game amongst the agents, where each individual's privacy if she participates, and therefore her participation choice, depends on the choices of the rest of the population.

We investigate symmetric Bayes-Nash equilibria, which in this game consist of *threshold strategies*, where all agents whose privacy requirements are weaker than a certain threshold participate and the remaining agents do not. We characterize these equilibria, which depend both on the noise announced by the analyst and the population size; present results on existence, uniqueness, and multiplicity; and discuss a number of surprising properties they display.

1. INTRODUCTION

In today's data-driven world, we encounter an increasing number of settings where substantial and varied information about individuals—from contact information, to purchase histories, to emails, financial data, medical data, opinions, locations, actions, and beyond—is aggregated and used. The Internet has made such data increasingly easier to obtain, and advances in databases and machine learning are rapidly improving the usefulness of this avalanche of data, allowing for better predictions, recommendations, and advertising (to name just a few of the many applications). Simultaneously, completely new uses for data— such as mobile health applications, or applications for well-being based on social data— are being invented using emerging technologies. From the perspective of an individual, these advances are a mixed blessing— one can certainly derive benefits from improved information aggregation, for example in the form of more accurate polling data or better suggestions of what movie to watch next. But there also are risks in sharing one's data: from malicious or accidental breaches, from poor anonymization techniques and other leaks, from unwanted inferences, and—perhaps most importantly—from threats that have not yet been discovered or imagined, leading each person to have their own level of discomfort or disutility from sharing their data.

There is by now a vast and rich literature on database privacy, particularly within the powerful framework of differential privacy. The main setting typically considered there is one where an analyst has access to a database of individuals' information and wishes to perform a computation with specific properties. Informally, the goal of the analyst is usually to optimize the quality of the output of the computation (according to some specific quality metric) while guaranteeing the individuals in the database a particular level of differential privacy. Recent work more carefully considers the incentives of the individuals in such a computation, proposing a variety of utility models that incorporate privacy considerations [Ghosh and Roth 2011; Xiao 2013; Chen et al. 2011; Nissim et al. 2012a; Ligett and Roth 2012], and laying the groundwork for the field of privacy-preserving mechanism design [McSherry and Talwar 2007; Xiao 2013; Chen et al. 2011; Nissim et al. 2012b,a; Huang and Kannan 2012; Kearns et al. 2012].

The realities of sharing personal data, however, have not kept pace with the advances in the formal study of privacy. Real-world situations involving the use of personal data typically say (at best) how an individual's data might be used, rather than promise a privacy guarantee, and an individual's option is (at best) to decide whether or not to agree to part with her data, rather than negotiate compensation for its use. In fact, a typical user might not even want— or be able to— specify such a compensation, although the typical user is likely to be able to say whether or not she would

be willing to share her data in any given scenario. Also, in many cases, one might argue, privacy comes not from a formal guarantee, but from a sense that one is likely safer 'hiding' in a larger crowd. In this way, one can view the decision of whether to participate in a computation emerge as an equilibrium between the potential participants—if others participate, perhaps I will too, and not much will be learnt about any of us. In this way, both privacy guarantees and participation decisions (and thus the database of participants itself) become an *endogenous* aspect of the model of computation. How should one think about privacy and accuracy in this new world of endogenous privacy and endogenous participation?

Our contributions. Our main contribution is a simple model where a privacy-sensitive population with privacy requirements decides whether or not to participate in a *pre-announced* noisy computation by an analyst, so that the database itself is *endogenously determined* by agents' participation choices. We model individuals as agents who each have some minimum privacy *requirement* that must be met for them to be willing to participate in any computation on their private data. An analyst announces a computation, along with the form and variance of the noise he will add to the outcome; this noise is some pre-decided quantity that is fixed prior to any participation decisions. The privacy an agent receives in this computation depends on both this noise *as well as* how many agents choose to participate in the database. Agents then make their participation decisions based on whether their expectation of how much privacy they will receive—given the announced computation and added noise, *and* their predictions about other agents' participation choices—meets their privacy requirement. (Note that there is no agent-specific compensation or payment to induce participation; agents participate simply if they are comfortable with the privacy they expect for their data, and do not participate otherwise). This gives rise to a *game* amongst the players, since whether an individual receives adequate privacy or not, and therefore participates or not, depends on the participation decisions of the remaining agents in the population. This model, of proposing a *computation* rather than a specific privacy guarantee, seems to us quite salient in that most real-world computations on sensitive data do not come with explicit formal privacy guarantees; also, by assuming only that individuals make a single binary decision about their privacy, we make minimal, fairly robust assumptions on how individuals reason about privacy.

Despite our focus on simplicity, to understand the nature of the trade-offs that can arise in such situations we necessarily must choose a concrete formal framework for privacy, and assume that agents can reason strategically to arrive at equilibrium outcomes. We analyze this privacy coordination game in a model where each individual has a *differential privacy* requirement r_i that is drawn from a commonly known distribution of privacy requirements $F(r)$, and each agent makes her participation decision based on her expected privacy when the remaining agents, with privacy requirements drawn from F, decide whether or not to participate according to some (known) strategy. We investigate symmetric Bayes-Nash equilibria in this game which consist of *threshold strategies*, where all agents whose privacy requirements are weaker than a certain threshold participate and the remaining agents do not. Our model and equilibrium analysis raise a number of subtleties that arise when computing on databases with endogenous participation, even when there are no correlations between individuals' privacy requirements and their data. First, an analyst might benefit from offering to add more noise than some 'bare minimum'—while adding higher noise has the direct effect of increasing inaccuracy, it could also potentially lead to a much larger subset of the population participating in the database, resulting in a final outcome that actually better approximates the true function of the population. Second, while it is natural to expect that adding more noise will lead to higher participation, this may not be the case, and equilibrium participation can potentially worsen with added noise. This happens precisely due to the presence of *multiple equilibria*, some of which improve with added noise as expected, whereas others become worse. However, the equilibrium corresponding to the *highest* participation improves with added noise, for any population size; we also demonstrate the existence of, and identify, a range of noise variances for which there is a unique equilibrium for any fixed population size. Finally, we show that, roughly speaking, equilibrium outcomes improve and converge towards full participation with diverging population size for

a reasonable range of noise parameters. We conclude in §4 by discussing how an analyst might use information about participation to improve performance and address some of the issues that arise from equilibrium multiplicity and the Bayesian information model.

1.1. Related work

An exciting literature has emerged on a variety of topics at the intersection of privacy, game theory and mechanism design. One fruitful line of inquiry has considered privacy-preserving mechanism design [McSherry and Talwar 2007; Nissim et al. 2012b; Xiao 2013; Chen et al. 2011; Huang and Kannan 2012; Nissim et al. 2012a], focusing on approximately- and exactly-truthful differentially private mechanisms. Some of this work also raises very interesting considerations in modeling the costs to individuals of privacy losses (for those works and additional papers that treat the topic see, e.g., [Kleinberg et al. 2001; Ghosh and Roth 2011; Xiao 2013; Chen et al. 2011; Nissim et al. 2012a; Ligett and Roth 2012]).

Closest to the present work is the recent literature on modeling and designing mechanisms for computation in settings where truthful reporting of *data* is not a concern, but where the data is sensitive and privacy losses incur a cost that must be appropriately compensated, at a rate that is private information and that agents can lie about [Ghosh and Roth 2011; Fleischer and Lyu 2012; Roth and Schoenebeck 2012; Ligett and Roth 2012; Dandekar et al. 2012]. The present paper differs from previous work in two main ways. First, our analyst simply announces a noisy computation, and does not attempt to elicit any costs from agents; there are no sophisticated mechanisms, analyst-supplied privacy guarantees, or individual-specific compensation for privacy costs. In the face of the announced noisy computation, each individual makes only a single binary participation decision. Second, as a result of the computation announced, a coordination game naturally emerges, allowing us to isolate the phenomenon of hiding-in-a-crowd that is naturally suggested by privacy in a game-theoretic model, and analyze its equilibrium behavior.

There has also been work on the design of markets and pricing schemes for personal data (rather than for privacy-preserving computation on that data (see, e.g., [Laudon 1993; Aperjis and Huberman 2012; Gkatzelis et al. 2012; Li et al. 2012]); on private equilibrium release [Kearns et al. 2012]; on a variety of approaches to running auctions in order to prevent unnecessary information about bids from being leaked (see, e.g., [Naor et al. 1999; Brandt and Sandholm 2008; Feigenbaum et al. 2010]); and on various other economic aspects of privacy (see, e.g., [Isaacman et al. 2011; Feigenbaum et al. 2012]).

The setting we study is clearly closely related to games studied in the economics literature, including public goods games (cf. [Varian 1994]), coordination games, and externalities;[1] despite these commonalities, there are significant differences between previous work in these areas and our model, arising due to the specific problem we study. For example, although the private computation in our setting has some characteristics of a public good, our concern here is not just whether *anyone* chooses to participate (analogously, whether the public good is funded), but *how many people* participate; similarly, individuals in our setting base their participation decisions on others' expected participation and do not receive assurance contracts [Bagnoli and Lipman 1989; Edlin et al. 1998] (and hence do risk receiving less privacy than expected; see §3.6 for discussion).

2. MODEL

There is a population of N agents, each of whom holds data of interest to an analyst. These N agents are the *potential* participants in the database that the analyst would like to perform a computation on, to learn some statistic about the population—the analyst would like his computation to return an answer close to that of the true computation on the full set of data from the entire population of size N. The analyst does not already hold the data from the population—agents will *choose* whether or not to contribute their data to the database *if* they are 'satisfied' with the privacy they expect to

[1] See, e.g., [Fudenberg and Tirole 1991; Leyton-Brown and Shoham 2008] or any standard text on game theory for discussion these subjects.

receive. We let n denote the number of *actual* participants in the database; naturally, $n \leq N$, the number of potential participants.

The analyst announces the computation he will perform on the participating agents' data, where by announcing the computation we mean that the analyst specifies what function will be applied to the data collected from individuals, as well as some pre-specified modification that will be made to the output to ensure (some degree of) privacy to agents. We suppose that this modification is *independent* of the actual number of participants, n—for example, the analyst might specify that he will add a particular form of random noise with variance μ where μ is some fixed quantity (dependent, perhaps, on the range in which the data is known to lie).

Given this announced computation, the actual privacy that each agent receives will depend on both the announced computation (including the specified modification), as well as typically the *number* of other agents who participate in the computation. The possible dependence on the number of agents arises because the analyst is interested in a function whose output does not scale with n (such as the mean), and the modification (such as the amount of noise added) is fixed independent of the number of actual agents who participate; thus, if a larger number of agents participate, each agent receives greater privacy.

Our agents are simple—they each have some minimum privacy requirement, and will participate in the computation if the privacy they expect to receive satisfies their requirement, and not if it doesn't (in addition to being a natural model for participation decisions, such threshold-based decisions arise naturally from existing utility models of privacy in the framework of differential privacy, as we will discuss shortly). Agents make participation decisions based on the announced computation and their own privacy requirement, leading to some realized database with $n \leq N$ participants. The analyst then performs the promised computation on the subset of the population who contributed to the database. The analyst's utility from the outcome of this computation depends on (i) the modification (such as the degree of noise perturbation) in the announced computation, and (ii) the size of the subset of the population that chose to participate in the computation. Intuitively, there is a potential trade-off here—a small announced modification will lead directly to less inaccuracy in the computation, but could also induce much lower participation and lead indirectly to higher inaccuracy, due to a small sample size, with respect to the true function estimated from the entire population N (which is the analyst's ideal).

Alternative modeling choices. One might naturally wonder about alternatives to our model of a pre-announced computation, and of privacy requirements. A first alternative modeling choice might be to announce a privacy *guarantee* itself, rather than announce a computation and perturbation. We do not choose this model for two reasons: (i) An analyst might not be willing or usefully able to perturb the output to the extent required to deliver on this privacy guarantee if participation is too low, and (ii) a typical user might be more likely to be able to decide whether or not she wants to participate in a given announced computation, than to parse or evaluate, a given guarantee. (We note here that analyzing outcomes in this setting will necessarily require that we use a formal framework, and that we assume users can reason rationally about their privacy requirements; however, within these constraints we adopt the simplest model of user behavior and privacy requirements possible).

A second alternative is one where an agent's privacy requirement is specified simply in terms of the number of other agents she needs in the database to be willing to participate, to capture the game between agents and the notion of hiding in a crowd. However, as we will see shortly, this number *depends* on the computation and the extent of the noise announced by the analyst—that is, two different computations and levels of added noise will lead to different degrees of privacy even with the same number of fellow participants. So any such specification must be a function of the added perturbation as well, at which point the two models become equivalent.

2.1. A mathematical model

We now set up a formal mathematical model for this setting, where for concreteness we need to make modeling choices regarding agent privacy requirements, the computation the analyst is interested in, and the noise he adds.

Privacy requirements. Our first modeling choice is to use the framework of *differential privacy* as agents' notion of privacy. (A brief overview of differential privacy and noise addition to achieve differential privacy, as well as a discussion of the specific nature of the privacy guarantees in our model, is in §2.3.) An agent's privacy *requirement*, given the differential privacy framework, is some real number greater than or equal to zero. (If desired, such threshold-like privacy requirements can be interpreted naturally as the choice of expected-utility maximizing agents in existing models for privacy costs in the literature; see §2.2.)

An agent with privacy requirement r_i participates in an announced noisy computation if her expected differential privacy guarantee is $p \leq r_i$ (for concreteness, we break ties in favor of participation—*i.e.*, an agent participates even if $p = r_i$). We assume that the privacy requirements of agents all lie in some interval $[r_{\min}, r_{\max}]$, where $r_{\min} \geq 0$ (we allow $r_{\max} = \infty$, in which case the interval is $[r_{\min}, \infty)$) and denote agent i's privacy requirement by r_i. We will suppose that each agent's privacy requirement is randomly drawn from a distribution $F(r)$ with support on $[r_{\min}, r_{\max}]$; we will assume for our analysis that F is continuously differentiable and strictly increasing on its support. Note that this distribution F is a distribution on privacy requirements and *not* on the data individuals hold, which is immaterial in our analysis.

Analyst's desired computation and added noise. We model each individual's private data as a real number $\in [0, 1]$. We suppose the computation of interest to the analyst is a low-sensitivity real-vector-valued computation on the population, such as a mean, or normalized histogram, computation. For such computations, since the impact of any one individual on the output decreases as the number of participants increases, adding Laplace noise with fixed standard deviation $\sim \frac{1}{\varepsilon}$ to the true value of the computation on the participants gives a differential privacy guarantee that improves with the total number of participants n as ε/n.[2] Informally, a larger value of ε corresponds to a lower level of added noise, and therefore less inaccuracy but also less privacy to participants.

Therefore, each agent makes her participation decisions based on whether she expects ε/n to be smaller than r_i or not; we describe agents' participation choices, information model, and the corresponding game and solution concept in §3. (The fact that a game arises does not depend either on our modeling choices of a differential privacy framework or addition of Laplace noise—it arises for *any* notion of privacy and privacy-inducing modifications as long as the modification announced is fixed and independent of the elicited participation; however, the specific structure of the game will depend on the formulation used.)

Remark. In this paper, we will assume that agents' privacy requirements r_i are *independent* of the actual data they hold, so that the accuracy of the final computation depends only on the *number*, rather than the specific subset, of agents that choose to participate. This assumption is not without loss of generality, and indeed there are settings (such as those involving sensitive medical data) where an agent's desire for privacy could well be correlated with her data. However, there are also other settings, such as collecting data on the Internet regarding (for example) browsing or shopping behavior, mobile applications for social health, or polls in online communities, where no particular data (or choice or behavior) is more or less incriminating or embarrassing than another—the differences in desired privacy levels arise simply because different people have different degrees of discomfort with sharing or revealing information about themselves, *whatever* that information.

2.2. Deriving requirement thresholds from a utility model with privacy costs

In this section, we show how threshold-like privacy requirements—making participation decisions by comparing expected privacy against a requirement r_i—arise directly when rational, expected-utility maximizing, agents make participation decisions in a simple utility model with privacy costs (as in prior literature, see e.g. [Ghosh and Roth 2011; Fleischer and Lyu 2012]). We note, though, that none of our results rely on this interpretation: our results only use the fact that each user has some minimum privacy requirement, no matter what its origin.

[2] See § 2.3 for discussion of Laplace noise and low-sensitivity computations.

Suppose an agent has a value v for the output of the computation, and a *privacy cost* that scales with the privacy guarantee that she receives from the computation. Suppose, for simplicity, that this privacy cost is linear in the level of differential privacy ε as in, e.g., [Ghosh and Roth 2011]: $c(\varepsilon) = c\varepsilon$. Then, the agent obtains utility $v - c\varepsilon$ when she participates in the computation and receives privacy ε. That is, her utility is nonnegative whenever

$$v - c\varepsilon \geq 0, \text{ or } \varepsilon < \frac{v}{c} = r.$$

Also, suppose there is a distribution on outcomes, with different numbers of participants, resulting in a distribution on the privacy guarantees received in each of those outcomes. The expected utility to the agent is

$$E[u] = E[v - c\varepsilon] = v - cE[\varepsilon],$$

so that

$$E[u] \geq 0 \Leftrightarrow E[\varepsilon] \leq r.$$

Thus, an agent will receive nonnegative expected utility if and only if her *expected privacy* guarantee $E[\varepsilon] \leq r$.

Suppose that this benefit from the computation only accrues to participants (as, for instance in applications that use data to provide personalized recommendations or outputs) when they have contributed data to the computation. Then, the utility from nonparticipation is zero, since the benefit and cost are both zero (since an agent who does not contribute her data receives perfect differential privacy, *i.e.*, $\varepsilon = 0$ (see §2.3)).

Therefore, an agent who is only making participation decisions will choose to participate if her expected privacy from participation is less than her threshold r_i, and not otherwise:

$$E[\varepsilon] \leq r_i.$$

Remark. We focus in this paper on mechanisms which only allow agents to make participation decisions—that is, we do not attempt to elicit agents' costs for privacy, or even just their thresholds. Rather, we assume that agents simply make the decision of whether to participate or not—while this simplifies analysis, a more significant reason is that a typical person is unlikely to be able to state a quantitative privacy requirement if asked, but is likely to also be able to say whether or not she is willing to share her data in any given scenario.

2.3. Modeling privacy

We consider preserving the privacy of individuals' data (we do not suppose that their privacy requirements are sensitive, as these are uncorrelated with their data), and so we think of an individual i's row in a database—corresponding to her data d_i—as her piece of the input to a mechanism that preserves differential privacy. We say that two databases D, D' are neighboring if they differ in at most one row. A guarantee of ε-differential privacy [Dinur and Nissim 2003; Dwork and Nissim 2004; Dwork et al. 2006] for a mechanism M requires that for all neighboring pairs of databases D, D' and for any possible outcome a of the mechanism, the following holds:

$$\Pr[M(D) = a] \leq \exp(\varepsilon) \Pr[M(D') = a].$$

The *sensitivity* of a function is the maximum by which the value of that function can change under any two neighboring databases. The mean and the normalized histogram have sensitivities $1/n$ and $2/n$, respectively, where we recall that n denotes the number of actual participants. The scaled symmetric exponential distribution with standard deviation $\sqrt{2}/\varepsilon$, denoted Laplace$(1/\varepsilon)$, has mass at x proportional to $\exp(-|x|\varepsilon)$. By the properties of Laplace noise [Dwork et al. 2006], the noisy computation announced by the analyst, wherein she adds noise $\sim$ Laplace$(1/\varepsilon)$ to a computation with sensitivity $\sim 1/n$, preserves $\frac{\varepsilon}{n}$-differential privacy.

We note here that the differential privacy parameter ε is indeed an imperfect metric by which to model privacy costs. In particular, since differential privacy provides a worst-case guarantee (over all

possible other members of the database), differential privacy is perhaps most useful as a *pessimistic* bound on privacy costs—under reasonable models, an individual's actual disutility from privacy loss when a differentially private computation is run on a specific input database may be less than that predicted by differential privacy. The threshold model we adopt helps sidestep this concern, since each individual's threshold may reflect her *interpretation* of the expected level of privacy provided by the announced additive noise.

We emphasize that differential privacy is merely one possible privacy framework and we choose it for concreteness; the idea of databases with endogenous participation could be applied more generally to other kinds of "hiding in a crowd" requirements—for instance, agents may simply need some minimum number of other agents contributing their data to feel comfortable about contributing their own.

3. EQUILIBRIUM ANALYSIS

We now analyze the nature of equilibrium participation when all N potential participants have common prior beliefs about the privacy requirements of the population, but only know their own actual privacy requirement with certainty. Specifically, every agent i knows the distribution $F(r)$ from which all other agents' privacy requirements are drawn, but not their actual draws; she observes her own draw r_i. (Again, we emphasize that this distribution F is a distribution on privacy requirements and *not* on the data individuals hold, which do not figure anywhere in our analysis.) Recall that the distribution F has support on $[r_{\min}, r_{\max}]$, and is assumed strictly increasing and continuously differentiable on its support, with bounded density $f(r)$.

Agent i makes her decision about whether or not to participate in the computation based on her own draw r_i from F, knowledge of F and N (the total number of potential participants),[3] and her beliefs about other agents' participation choices, as follows. Suppose all other agents use a strategy $s(r)$ to decide whether or not they will participate as a function of their privacy requirement r. Agent i participates if her expected differential-privacy guarantee $p(s(r), N, \varepsilon)$, meets or is better than her own privacy requirement r_i, where the expectation is taken over agents' random draws of r_j from F.

For a given set of $N-1$ iid draws from $F(r)$, let $n(s(r))$ denote the number of these remaining $N-1$ agents who participate under the strategy $s(r)$. Then, agent i participates if her expected differential privacy guarantee from participating, $p(s(r), N, \varepsilon) = E[\frac{\varepsilon}{n(s(r))+1}]$,[4] satisfies $p \leq r_i$, and does not participate otherwise (recall that we assume throughout that an agent who is indifferent between participating and not breaks ties in favor of participation.)

Symmetric Bayes-Nash equilibria. We will focus on symmetric strategies. A strategy $s(r^*)$ constitutes a symmetric Bayes-Nash equilibrium if no agent can improve her expected utility by deviating from $s(r)$ and making her participation decisions according to a different strategy $\hat{s}(r)$, when the remaining $N-1$ agents all make their participation decisions according to $s(r)$. [5]

3.1. Characterizing symmetric equilibria

We first define threshold strategies.

Definition 3.1. A *threshold* strategy $s(r^*)$ is a strategy where an agent i participates if her privacy requirement r_i is greater than or equal to the threshold r^*, and does not participate otherwise.

---

[3] The assumption that agents know N is justifiable in some settings and less so in others, and we leave open the problem of exploring uncertainty or unknown N as a direction for further exploration

[4] Here, the expectation is taken over the $(N-1)$ other players' draws from F.

[5] We choose a Bayes-Nash equilibrium analysis since it is more reasonable to expect that individuals have some idea about the nature of the population's privacy sensitivities rather than to expect that they actually know other agents' privacy requirements as in a full-information setting, especially since an individual typically will not even know the identities of the other players in the game. A setting where the analyst reveals other agents' participation decisions is another natural choice for the information model, and is one we discuss briefly in §4.2.

We will, *without loss of generality*, restrict our agents to using threshold strategies. This is because agent i's unique best response for any given strategies of other agents is a threshold strategy: consider the threshold $r_i^*(s(r)) = p(s(r), N, \varepsilon) = E[\frac{\varepsilon}{n(s(r))+1}]$. If agent i participates when she has requirement r, by definition $r \geq p(s(r), N, \varepsilon) = r_i^*(s(r))$, and conversely an agent does not participate if $r < p(s(r), N, \varepsilon)$. So agent i participates iff $r \geq r_i^*(s(r))$, which is a threshold strategy with threshold $r_i^*(s(r))$.

A threshold strategy *equilibrium* is one where all agents whose privacy requirements are above a certain threshold r^* participate and all agents with stricter privacy requirements, *i.e.*, with requirements below r^*, do not participate, and no agent can benefit by deviating from this strategy. We have the following easy lemma.

LEMMA 3.2. *A threshold strategy $s(r^*)$, $r^* \in (r_{\min}, r_{\max})$, is a symmetric Bayes-Nash equilibrium iff $E[\frac{\varepsilon}{n(r^*)+1}] = r^*$, where $n(r^*)$ is the binomial random variable denoting the number of the $N-1$ agents who draw values $r_i \geq r^*$.*

PROOF. An agent should participate iff $r_i \geq p(s(r^*), N, \varepsilon) = E[\frac{\varepsilon}{n(r^*)+1}]$, where the last term is the expected privacy to the agent when $N-1$ other agents participate according to threshold r^*. For r^* to be an equilibrium, every agent with $r_i \geq r^*$ must obtain expected privacy weakly better than her requirement, *i.e.*, less than or equal to r_i, and every agent with $r_i < r^*$ must expect to receive privacy strictly worse (*i.e.*, greater) than r_i from participating. So r^* is an equilibrium threshold iff $r^* = E[\frac{\varepsilon}{n(r^*)+1}]$. □

We abuse notation to let $p(r, N, \varepsilon)$ denote the expected privacy to an agent from participating when all other agents use a threshold strategy $s(r)$. The next result characterizes equilibrium thresholds.

THEOREM 3.3. *For any given N and ε, and a distribution F:*

(1) $r_{\min}$ is an equilibrium threshold iff $\frac{\varepsilon}{N} \leq r_{\min}$.
(2) A value $r^* \in (r_{\min}, r_{\max}]$ is an equilibrium threshold if and only if it satisfies the following condition:

$$\frac{\varepsilon}{N} \frac{1 - F^N(r^*)}{1 - F(r^*)} = r^*. \tag{1}$$

PROOF. The first statement follows immediately since at $r^* = r_{\min}$, all agents participate so $p(r_{\min}, N, \varepsilon) = \frac{\varepsilon}{N}$, and this is an equilibrium iff $\frac{\varepsilon}{N} \leq r$ for all $r \in [r_{\min}, r_{\max}]$.

For the second statement we evaluate $p(r, N, \varepsilon)$. The differential privacy when n agents participate is $\frac{\varepsilon}{n}$, and each of the remaining $N-1$ agents participates with probability $1 - F(r)$ when the threshold is r. Noting that $1 - F(r) > 0$ for $r \in (r_{\min}, r_{\max})$ (since F is strictly increasing, $F(r) < F(r_{\max}) = 1$), we have

$$\begin{aligned}
\mathbb{E}\left[\frac{\varepsilon}{n}\right] &= \varepsilon \sum_{k=0}^{N-1} \frac{1}{k+1} \binom{N}{k} (1 - F(r))^k F(r)^{N-1-k} \\
&= \frac{\varepsilon}{N(1 - F(r))} \sum_{k=0}^{N-1} \binom{N}{k+1} (1 - F(r))^k F(r)^{N-1-k} \\
&= \frac{\varepsilon}{N(1 - F(r))} (1 - F^N(r)).
\end{aligned}$$

From Lemma 3.2, r^* is an equilibrium threshold iff this expectation equals r^*, proving the result. □

The function on the left-hand side in Theorem 3.3 appears repeatedly in our equilibrium analysis:

Definition 3.4 ($p(r, N, \varepsilon)$). The expected differential privacy guarantee with added Laplace$(1/\varepsilon)$ noise to a participating agent, when $N-1$ other agents with privacy require-

ments drawn from F participate if and only if their draws r_i exceeds a threshold r, is

$$p(r, N, \varepsilon) = \frac{\varepsilon}{N} \frac{1 - F^N(r)}{1 - F(r)}.$$

We will sometimes drop the arguments N, ε when the values of N, ε are clear from context.

The following two propositions lead to monotonicity properties of $p(r, N, \varepsilon)$ that are used repeatedly in our equilibrium analysis.

PROPOSITION 3.5. *The function $\frac{1-x^N}{N}$ strictly decreases with increasing N for any $x \in (0, 1)$.*

PROPOSITION 3.6. *The function $\frac{1-F^N(r)}{1-F(r)}$ is increasing in r for $r \in [r_{\min}, r_{\max}]$ and any $N > 0$.*

3.2. Equilibrium existence

Having understood what values of r can constitute a threshold equilibrium, the first question we address is existence of equilibrium. We will focus throughout on equilibria with non-trivial participation, *i.e.*, with thresholds $r^* < r_{\max}$— the expected participation at a threshold equilibrium r^* is $\mathbb{E}[n] = N(1-F(r^*))$, which is zero at $r^* = r_{\max}$. How does existence of equilibria with non-trivial participation depend on the value of added noise ε, and population size N?

THEOREM 3.7. *Consider any pair (N, ε). An equilibrium exists if $\varepsilon/N \leq r_{\min}$ (at $r^* = r_{\min}$). If $\varepsilon/N > r_{\min}$, then*

(1) There exists at least one equilibrium $r^* < r_{\max}$ for all $\varepsilon < r_{\max}$.
(2) For all $\varepsilon \geq r_{\max}$, there is either no equilibrium $r^* < r_{\max}$, or multiple such equilibria r^*, except on a set of ε of measure zero.

PROOF. Recall that F is continuous, and note that the function

$$p(r) = \frac{\varepsilon}{N} \frac{1 - F^N(r)}{1 - F(r)} = \frac{\varepsilon}{N}(1 + F(r) + \ldots + F^{N-1}(r))$$

is therefore continuous in r for all $r \in [r_{\min}, r_{\max}]$.

The first statement follows immediately from Theorem 3.3. For the second, suppose $\varepsilon/N > r_{\min}$. An equilibrium exists at any value of $r \in (r_{\min}, r_{\max}]$ at which $p(r, \varepsilon) = r$. At $r = r_{\min}$, $p(r, \varepsilon) = \frac{\varepsilon}{N} > r_{\min}$ by assumption. At $r = r_{\max}$, $p(r, \varepsilon) = \varepsilon$, which is strictly smaller than $r_{\max}$ if $\varepsilon < r_{\max}$. Therefore, $p(r, \varepsilon) - r > 0$ at $r = r_{\min}$ and $p(r, \varepsilon) - r < 0$ at $r = r_{\max}$. Since $p(r, \varepsilon) - r$ is a continuous function of r on $[r_{\min}, r_{\max}]$, and is positive at $r_{\min}$ and negative at $r_{\max}$, it must cross 0 at least once in the interval $[r_{\min}, r_{\max}]$ by the intermediate value theorem.

If $\varepsilon > r_{\max}$, $p(r, \varepsilon) - r$ is positive at $r = r_{\min}$ and $r = r_{\max}$. So either (i) there is no value of r at which $p(r, \varepsilon) = r$, or (ii) if $p(r, \varepsilon) - r$ *crosses* 0 at some r to become negative, it must cross 0 again at some $r' > r$ since it must eventually be positive at $r = r_{\max}$. The only situation where there is a unique equilibrium is when $p(r, \varepsilon) = r$ for a unique $r \in (r_{\min}, r_{\max})$, *and* $\frac{\partial p(r,\varepsilon)}{\partial r} = 1$. But this can only happen for exactly only value of ε because of the linear homogeneous dependence of $p(r, \varepsilon)$ on ε. □

We note here that when $\varepsilon < r_{\max}$, there is an odd number of equilibria, and when $\varepsilon \geq r_{\max}$, there is an even number (including possibly 0) of equilibria, except on a measure zero set of ε (where the function $p(r, \varepsilon)$ is tangent to r at $p(r) - r = 0$, leading to equilibria without a zero crossing.)

The first question an analyst working with a population of size N might ask is at which values of ε (the noise he offers to add) do equilibria exist. Our next theorem says that for any N, there is an *interval* of ε values at which equilibria exist; the size of this interval grows with N.

THEOREM 3.8. *For any fixed N, the set of ε values for which an equilibrium with $r^* < r_{\max}$ exists is an interval with endpoints 0 and $\varepsilon_e(N)$, where $\varepsilon_e(N) \geq r_{\max}$. (The interval is semi-open and does not contain the right endpoint $\varepsilon_e(N)$ if $\varepsilon_e(N) = r_{\max}$, and is closed if $\varepsilon_e(N) > r_{\max}$). Further, $\varepsilon_e(N)$ is increasing in N.*

PROOF. We first argue that this set of ε values is an interval. Let ε_1 be such that an equilibrium $r^* < r_{\max}$ exists, and consider any $\varepsilon_2 < \varepsilon_1$. First, if $r^*(\varepsilon_1) = r_{\min}$, $r_{\min} \geq p(r_{\min}, \varepsilon_1) = \frac{\varepsilon_1}{N} > \frac{\varepsilon_2}{N}$, so $r_{\min}$ is also an equilibrium for ε_2. So suppose $p(r_{\min}, \varepsilon_1) - r_{\min} > 0$, and $p(r_{\min}, \varepsilon_2) - r_{\min} > 0$ also (since otherwise $r^*(\varepsilon_2) = r_{\min}$ is an equilibrium). Note that since an equilibrium exists for ε_1, $r \geq p(r, \varepsilon_1)$ for some $r \in (r_{\min}, r_{\max})$. But then $p(r, \varepsilon_2) < 0$ at this r since $p(r, \varepsilon)$ is strictly decreasing in ε. So $p(r, \varepsilon_2) - r$ must cross zero by continuity, so an equilibrium exists for ε_2, and the set of ε for which an equilibrium exists is an interval.

Now, since an equilibrium with $r^* < r_{\max}$ exists (by continuity of $p(r)$) iff

$$\frac{\varepsilon}{N} \frac{1 - F^N(r)}{1 - F(r)} \leq r, \text{ or } \varepsilon \leq N \frac{r(1 - F(r))}{1 - F^N(r)}$$

for some $r \in [r_{\min}, r_{\max})$, let

$$\varepsilon_e(N) = \max_{r \in [r_{\min}, r_{\max}]} N \frac{r(1 - F(r))}{1 - F^N(r)}. \tag{2}$$

Then an equilibrium with $r^* < r_{\max}$ exists for all $\varepsilon \leq \varepsilon_e(N)$ if this maximum occurs at $r_0 < r_{\max}$, and for $\varepsilon < \varepsilon_e(N)$ if $r_0 = r_{\max}$; note that an equilibrium cannot exist for any ε greater than this value.

Finally, the endpoint $\varepsilon_e(N)$ defined by (2) is increasing in N by Proposition 3.5. □

3.3. Monotonicity

Recall that the analyst's utility depends on the noise ε via both a direct effect—from the inaccuracy introduced by adding Laplace$(1/\varepsilon)$ noise—and the extent of participation, n, elicited in equilibrium when the announced noise perturbation is ε; a larger sample n should lead to a more accurate estimate. We now consider a fixed population size N, and ask how expected equilibrium participation, which is determined by the equilibrium threshold r^*, behaves as a function of the offered noise ε.

Our results here illustrate a non-intuitive fact about equilibrium behavior in this setting. A smaller value of ε corresponds to more added noise—with a fixed number of participants, this should lead to greater privacy for agents. So it is reasonable to expect that decreasing ε should cause more agents to participate. However, as we will see, this need not always be the case—and in fact, it is usually not the case, unless there is a unique equilibrium threshold. This unexpected behavior occurs due to *multiplicity* of equilibria—while the *'best'* equilibrium (the equilibrium with the highest expected participation) improves with increasing added noise, there is typically another equilibrium threshold whose value increases when there are multiple equilibria, at which expected participation drops when a noisier computation is announced. We note that such multiple equilibria are not uncommon in settings with externalities or network effects ([Easley and Kleinberg 2010]), or in coordination games , and in fact it is typical, even in our specific setting, for multiple equilibria to exist (even with $\varepsilon < r_{\max}$ where Theorem 3.7 does not guarantee multiple equilibria—an easy instance of this can be seen when F is the exponential distribution and N is not too large).

We formalize this behavior of equilibrium thresholds in the next two results.

Definition 3.9 (*Best equilibrium r_b^**). We refer to the smallest value of $r \in [r_{\min}, r_{\max}]$ that is a threshold equilibrium as the 'best' equilibrium for given values of N, ε, and denote it by $r_b^* = r_b^*(N, \varepsilon)$ (note that this is the equilibrium with the highest expected participation). If there is no solution to (1) in $r \in [r_{\min}, r_{\max}]$, we define $r_b^* = \infty$.

THEOREM 3.10. *The best equilibrium threshold $r_b^*(N, \varepsilon)$ is non-decreasing in ε for each N.*

PROOF SKETCH. Consider again the function $p(r,\varepsilon)-r$, and $\varepsilon_1 < \varepsilon_2$. If $r_b^*(\varepsilon_2) = \infty$, *i.e.*, there is no r satisfying the conditions of Theorem 3.3 in $[r_{\min}, r_{\max}]$, the statement holds vacuously. If $r_b^*(\varepsilon_2) = r_{\min}$, then $r_b^*(\varepsilon_1) = r_{\min}$ also, since $p(r_{\min}, \varepsilon_1) < p(r_{\min}, \varepsilon_2) < r_{\min}$ so that $r_{\min}$ is an equilibrium from Theorem 3.3, and therefore the best equilibrium.

So suppose there is an equilibrium $r_{\min} < r_b^*(\varepsilon_2) \leq r_{\max}$, so that $p(r_b^*(\varepsilon_2), \varepsilon_2) - r_b^*(\varepsilon_2) = 0$. Since $\varepsilon_1 < \varepsilon_2$, $p(r_b^*(\varepsilon_2), \varepsilon_1) - r_b^*(\varepsilon_2) < 0$ (note that $p(r,\varepsilon)$ is strictly increasing in ε for fixed r). Now since $p(r,\varepsilon) - r$ is $\frac{\varepsilon}{N} > 0$ at $r = r_{\min}$, by continuity of p in r, there exists some $r \in (r_{\min}, r_b^*(\varepsilon_2))$ at which $p(r,\varepsilon_1) - r = 0$. Therefore, $r_b^*(\varepsilon_1) < r_b^*(\varepsilon_2)$. □

The possibility of worse equilibrium participation despite the increase in added noise arises precisely because of the presence of additional equilibria beyond r_b^*—as we just proved in Theorem 3.10, the *best* equilibrium r_b^* behaves as one would expect, with increasing amounts of added noise leading to a decreased threshold and correspondingly increasing expected participation $N(1 - F(r_b^*))$ (recall that we are considering a fixed population N). However, at additional equilibria with $r^* > r_b^*$, there can be an unexpected reversal of participation. In Theorem 3.11, we state a sharper result than simply saying that equilibria can sometimes get worse, saying precisely which equilibria have improving participation and which get worse; to state the result we first need some notation.

Fix N and for a given ε, let $r_i^*(\varepsilon)$ denote the i-th smallest equilibrium threshold less than $r_{\max}$, *i.e.*, the ith smallest solution to the equation $p(r,\varepsilon) = r$ in the interval $[r_{\min}, r_{\max})$. Then, we have the following result regarding the local behavior of the thresholds r_i^* with ε (recall that the function F is continuously differentiable, so $p(r) - r$ is differentiable as well on $[r_{\min}, r_{\max})$).

THEOREM 3.11. *Fix N and ε, and suppose all solutions r_i^* are* zero crossings *of $p(r) - r = 0$ for this N and ε. If the function $p(r) - r$ crosses 0 sloping upward at r_i^*, then adding more noise (a decrease in ε) worsens the equilibrium, and if it crosses 0 sloping downward at r_i^*, then adding more noise improves the equilibrium. That is,*

(1) If $\frac{\partial}{\partial r}(p(r,\varepsilon) - r) > 0$ at $r = r_i^*$, then $r_i^*(\varepsilon)$ is locally decreasing in ε.
(2) If $\frac{\partial}{\partial r}(p(r,\varepsilon) - r) < 0$ at $r = r_i^*$, then $r_i^*(\varepsilon)$ is locally increasing in ε.

PROOF SKETCH. First suppose that $\frac{\partial}{\partial r}(p(r,\varepsilon) - r) > 0$ at r_i^*. Then there exists δ small enough so that $p(r_i^* - \delta, \varepsilon) - (r_i^* - \delta) < 0$, and $p(r_i^* + \delta, \varepsilon) - (r_i^* + \delta) > 0$. Note that for $\varepsilon' < \varepsilon$, $p(r_i^*, \varepsilon') - (r_i^*) < 0$, since $p(r,\varepsilon)$ is increasing in ε for fixed r. Now, because $p(r,\varepsilon)$ is also continuous in ε, the value δ can be chosen to be adequately small so that $p(r_i^* + \delta, \varepsilon') - (r_2^* + \delta) > 0$ also. Then, we have that $p(r,\varepsilon') - r < 0$ at $r_i^*(\varepsilon)$ and $p(r,\varepsilon') - r < 0$ at $r = r_i^*(\varepsilon) + \delta$. So by continuity of $p(r) - r$ in r and the intermediate value theorem, there is a solution $r_i^*(\varepsilon')$ to $p(r,\varepsilon') - r$ in the interval $[r_i^*(\varepsilon), r_i^*(\varepsilon) + \delta]$, *i.e.*, the corresponding equilibrium threshold at ε' is larger than the equilibrium threshold $r_i^*(\varepsilon)$.

The argument can be reversed for the case when the function is sloping downward at an equilibrium threshold r_i^* to show that for such equilibria, decreasing ε causes the solution r_i^* to locally decrease, *i.e.*, equilibria corresponding to downward-sloping intersections of $p(r)-r$ with 0 improve with added noise. □

We require the caveat about zero crossings, as opposed to being tangent to the x-axis, because of the numbering issue with the equilibria—if there is a threshold r_i^* where $p(r) - r$ is tangent to 0, a threshold equilibrium near this r_i^* fails to exist for perturbations of ε in one direction, and causes the numbering of the solutions to change in a way that leads to a discontinuity in r_i^*.

Self-fulfilling expectations equilibria and stability. Suppose, rather than thinking of F as the CDF of the distribution from which agents' privacy requirements are randomly drawn, we imagine that $F(r)$ specifies the *exact fraction* of the population with privacy requirements below r (the analysis with this interpretation of F is not identical, but very similar to that we have presented so far). Then a solution r^* to the equation $p(r, N, \varepsilon) = r$ is a *self-fulfilling expectations equilibrium* [Easley

and Kleinberg 2010]: if the population believes that all individuals with privacy requirements weaker than (*i.e.*, above) r^* will participate, this is exactly the set of individuals that will participate. It can be shown that equilibria r^* where $\frac{\partial}{\partial r}(p(r,\varepsilon) - r) < 0$ correspond exactly to *stable* equilibria in the sense that a dynamic process, where agents repeatedly revise participation decisions based on realized privacy levels, converges to r^* starting from any r near r^*. Similarly, unstable equilibria are those where $\frac{\partial}{\partial r}(p(r,\varepsilon) - r) < 0$.

Our theorem says that the equilibria that 'behave well', *i.e.*, as expected in the sense of monotonicity with respect to ε, are precisely the set of stable equilibria and the equilibria that 'behave badly' are the unstable ones. We note here that this behavior is reminiscent of, and related to, the 'correspondence principle' [Samuelson 1983] regarding comparative statics and stability of equilibria in economics, although the specific structure of our model requires us to derive this behavior directly for our setting.

3.4. Uniqueness

We have just seen that for a given population N, multiple equilibria may exist for at least some values of ε. An analyst may prefer to use values of ε that do not support multiple equilibria (to avoid the low participation corresponding to equilibria much larger than r_b^*, and also because of the undesirable monotonicity properties at some of these equilibria), choosing to optimally pick ε from amongst the set of ε at which there is a unique equilibrium. We next investigate uniqueness of equilibria as a function of ε for a given population N.

What does the set of ε at which there is a unique equilibrium look like? Unlike the set of ε values at which an equilibrium *exists*, which is always an interval, the set of ε with unique equilibria need not be an interval for general distributions F. However, this set does always *contain* an interval, the upper endpoint $\varepsilon_u(N)$ of which tends to $r_{\max}$ as N diverges. This interval thus grows to essentially the largest possible set of ε with unique equilibria, since recall from Theorem 3.7 that there are always multiple equilibria (if any exist) for $\varepsilon > r_{\max}$ for any N.

THEOREM 3.12. *For any N, the set of ε values for which there is a unique equilibrium $r^* < r_{\max}$ contains an interval of the form $[0, \varepsilon_u(N)]$. Further, the upper bound of this interval converges to $r_{\max}$ as N diverges:* $\lim_{N\to\infty} \varepsilon_u(N) = r_{\max}$.

PROOF. Consider $\varepsilon \in [0, \varepsilon_e(N)]$ for which an equilibrium exists. For uniqueness, there should be no solution to $p(r) = r$ in the interval $(r_b^*(N), r_{\max})$, where recall that $r_b^*(N)$ is the best equilibrium, or smallest solution to this equation in $[r_{\min}, r_{\max})$.

Now, since $p(r)$ is increasing in r for any particular ε by Proposition 3.6, any solution to $p(r) = r$ for $r \in (r_b^*, r_{\max})$ must lie in the subinterval (r_b^*, ε), since $p(r) < p(r_{\max}) = \varepsilon$ for all $r \in (r_b^*, r_{\max})$.

Since $p(r_b^*(\varepsilon), \varepsilon) - r_b^*(\varepsilon) = 0$, a *sufficient* condition for there to be no further solution to $p(r,\varepsilon) - r = 0$ in this interval is if $p(r, N) - r$ is decreasing in r on $(r_b^*, \varepsilon]$; by the previous argument this means there is no additional equilibrium in $(r_b^*, r_{\max})$ leading to uniqueness. For $p(r) - r$ to be decreasing on this interval, we need the derivative of $p(r) - r$ to be negative for all $r \in (r_b^*, \varepsilon]$, *i.e.*,

$$h(r, N, \varepsilon) = \frac{\partial}{\partial r}(p(r,\varepsilon) - r) = \frac{\varepsilon f(r)}{1 - F(r)} \left[\frac{1 - F^N(r)}{N(1 - F(r))} - F^{N-1}(r) \right] - 1 < 0.$$

We now argue that for a particular N, the set of ε values satisfying this sufficient condition for uniqueness is an interval, so that the set of values of ε with unique equilibria contains such an interval. Note that if $h(r, \varepsilon_1) < 0$ for all $r \in [0, \varepsilon_1]$, $h(r, \varepsilon_2) < 0$ for all $r \in [0, \varepsilon_2]$ as well for any $\varepsilon_2 < \varepsilon_1$: h is strictly decreasing in ε so $h(r, \varepsilon_2) < h(r, \varepsilon_1)$ for all r, and we require $h(r, \varepsilon_2) < 1$ to hold on a subinterval of the interval $[0, \varepsilon_1]$ where $h(r, \varepsilon_1) < 1$. Therefore the set of ε satisfying this sufficient condition is an interval of the form $[0, \varepsilon_u(N)]$.

To argue that $\lim_{N\to\infty} \varepsilon_u(N) = r_{\max}$, we need to show that for any $\delta > 0$, there is N large enough such that $|r_{\max} - \varepsilon_u(N)| < \delta$. Consider any $\bar{\varepsilon} < r_{\max} - \delta$. From Theorem 3.16, we

know that there is $N(\bar{\varepsilon})$ large enough so that there is a unique equilibrium at $\bar{\varepsilon}$ for all $N \geq N(\bar{\varepsilon})$. Therefore for all $N \geq N(\bar{\varepsilon})$, $\varepsilon_u(N) \geq \bar{\varepsilon}$, since $\bar{\varepsilon}$ belongs to the interval $[0, \varepsilon_u(N)]$ for all such N. Finally, $\varepsilon(N) \leq r_{\max}$ for all N since there are either no equilibria or multiple equilibria for $\varepsilon \geq r_{\max}$ from Theorem 3.7.

Therefore $|r_{\max} - \varepsilon(N)| \leq |r_{\max} - \bar{\varepsilon}| < \delta$, as required. So $\lim_{N\to\infty} \varepsilon_u(N) = r_{\max}$. □

We note here that multiple equilibria exist for $\varepsilon > r_{\max}$ even as the population size diverges (that is, even a very large population cannot guarantee uniqueness for all values of ε)— this follows immediately from Theorem 3.7.

3.5. Equilibrium behavior with diverging N

Finally, we investigate how agent participation in equilibrium behaves as a function of the population size N. Is it always a good idea for the analyst to try to reach as large a potential population as possible—how does participation in the best equilibrium behave as the population size grows, and can we be sure to eventually obtain the high participation equilibrium for a large enough population, *i.e.*, how does the multiplicity of equilibria behave with diverging N? And, is it possible, for *any* distribution of privacy requirements F, to eventually obtain participation from the entire population of agents, for a large enough population?

We start with existence, where our first result says that for any ε the analyst may wish to choose, however large, there is eventually a population size at which a non-trivial equilibrium threshold (*i.e.*, with threshold strictly less than $r_{\max}$ and non-zero participation) exists.

THEOREM 3.13 (EXISTENCE). *Consider any distribution F and any value of ε. There exists $\hat{N} = \hat{N}(\varepsilon)$ such that there is a threshold equilibrium with $r^* < r_{\max}$ for all $N \geq \hat{N}$.*

PROOF. Such an equilibrium exists if $p(r, N, \varepsilon) - r = 0$ for some $r \in [r_{\min}, r_{\max})$. We drop the dependence on ε for the remainder of this section.

First, if $r_{\min} > 0$, choose N large enough so that $p(r_{\min}, N) = \varepsilon/N < r_{\min}$: if $p(r_{\min}, N) \leq r_{\min}$, then $r_{\min}$ is an equilibrium threshold from Theorem 3.3, and we are done.

So suppose $r_{\min} = 0$. Consider some $r_0 \in (r_{\min}, r_{\max})$, and note that $r_0 > r_{\min} = 0$, and $r_0 < r_{\max}$ so that $F(r_0) < 1$. Let

$$N_0 > \frac{\varepsilon}{r_0(1 - F(r_0))}.$$

Then, for all $N > N_0$,

$$p(r_0, N) = \frac{\varepsilon}{N} \frac{1 - F^N(r_0)}{1 - F(r_0)} < \frac{\varepsilon}{N} \frac{1}{1 - F(r_0)} < r_0.$$

Therefore $p(r_{\min}) > r_{\min}$, and $p(r_0) < r_0$ for all $N > N_0$, so by continuity of $p(r) - r$, there is at least one solution to $p(r) = r$ for $r \in (r_{\min}, r_0)$, leading to a non-trivial equilibrium threshold as desired. □

While a (non-trivial) equilibrium always exists for large enough N, this equilibrium need not be unique, as we have seen already. The following result, paralleling Theorem 3.10, tells us that even when there are multiple equilibria, the best equilibrium r_b^* (Definition 3.9)—corresponding to the maximum expected participation—improves as the population N grows larger for any fixed ε.

THEOREM 3.14 (MONOTONICITY). *Consider any distribution F and any fixed ε. The best equilibrium $r_b^*(N)$ decreases monotonically with N.*

PROOF. Let $\hat{N}(\varepsilon)$ be such that a non-trivial equilibrium $r^* < r_{\max}$ exists for all $N \geq \hat{N}$, and note that $r_b^*(N) \geq r_{\max}$ by definition for all smaller N, *i.e.*, $N < \hat{N}$. So we only need consider values of N that are greater than $\hat{N}$.

Consider $N_2 > N_1 \geq \hat{N}$, and let $r_b^*(N_1)$ and $r_b^*(N_2)$ denote the corresponding best equilibrium thresholds. From Proposition 3.5, $\frac{1-F^N(t)}{N}$ is decreasing in N, so $p(r,N) - r$ decreases as N increases for any *fixed* r, and specifically for $r = r_b^*(N_1)$:

$$p(r_b^*(N_1), N_2) - r_b^*(N_1) < p(r_b^*(N_1), N_1) - r_b^*(N_1) = 0.$$

Therefore, since $p(r, N_2) - r$ is ε/N_2 is strictly greater than 0 at $r = r_{\min}$ and strictly smaller than 0 at $r = r_b^*(N_1)$, $p(r, N_2) - r = 0$ for some $r \in [r_{\min}, r_b^*(N_1))$ by continuity of $p(r) - r$. So a solution, and therefore the smallest solution, of $p(r, N_2) - r = 0$ occurs in the semi-open interval $[r_{\min}, r_b^*(N_1))$, so that $r_b^*(N_2) < r_b^*(N_1)$. $\square$

In fact, the best equilibrium not only decreases with increasing N, but converges to $r_{\min}$ in the limit of a diverging population, as the next results shows—this means that in the limit as N diverges, almost everyone in the population participates in the computation in equilibrium.

THEOREM 3.15 (LIMITING BEHAVIOR OF r_b^*). *Fix ε. For any continuous F, the best equilibrium threshold r_b^* converges to $r_{\min}$ in the limit as $N \to \infty$; the expected number of participants in this equilibrium diverges as $\Theta(N)$.*

PROOF. From Theorem 3.14, the sequence $r_b^*(N)$ is decreasing, and is bounded below by $r_{\min}$. Therefore the sequence $r_b^*(N)$ has a limit. Suppose this limit is distinct from $r_{\min}$, say $r_b^*(N) \to c > r_{\min} \geq 0$. with $F(c) > 0$.

The equation (1) satisfied by any equilibrium r^*, and therefore r_b^* can be rearranged to:

$$\frac{r^*(1 - F(r^*))}{1 - F^N(r^*)} = \frac{\varepsilon}{N}$$

if $r^* > r_{\min}$. (Since we assumed that r_b^* converges down to $c > r_{\min}$, this is the equation that must be satisfied by $r_b^*(N)$ for all N, and the inequality that must be satisfied for $r_b^* = r_{\min}$ does not apply.) Since F is continuous,

$$\lim_{N\to\infty} \frac{r_b^*(N)(1 - F(r_b^*(N)))}{1 - F^N(r_b^*(N))} = \frac{c(1 - F(c))}{1 - F^N(c)} > 0,$$

whereas $\lim_{N\to\infty} \frac{\varepsilon}{N} = 0$, a contradiction. Therefore we must have $\lim_{N\to\infty} r_b^*(N) \to r_{\min}$. The statement about expected participation follows immediately because $\mathbb{E}[n] = N(1 - F(r_b^*))$ and $F(r_b^*) \to 0$ as $r_b^* \to r_{\min}$ since F is continuous and $F(r_{\min}) = 0$. $\square$

As we observed when varying ε for fixed N, the multiplicity of equilibria again means that while one would expect that equilibria should improve with increasing population size N, in fact not all equilibria improve with N. The behavior of the equilibrium thresholds with N essentially follows the same pattern in Theorem 3.11— for equilibria that arise from solutions to $p(r, N) - r = 0$ where $p(r) - r$ has a downward-sloping zero-crossing, increasing N improves the equilibrium threshold and participation; whereas equilibria where $p(r) - r$ has an upward-sloping zero-crossing become worse with increasing N, displaying behavior contrary to the expected monotonicity. Stating the result formally for N is tricky because N is an integer, so we do not have continuity of $p(r, N)$ in N—roughly, there might be no intersection with 0 as we move discretely from N to $N-1$ or $N+1$; hence, we omit a formal statement of the result. Again, this unexpected behavior due to multiplicity of equilibria leads immediately to the question about uniqueness.

Our final result says that there is eventually a unique equilibrium for any given value of $\varepsilon < r_{\max}$ (so that, for example, an analyst who has decided on some fixed ε can still hope for a unique equilibrium by attempting to reach a large population N); this solitary equilibrium is the one with the nice properties described earlier. We note here that this condition on ε is essentially the best possible, since from Theorem 3.7 there are always multiple equilibria (if any exist) for $\varepsilon \geq r_{\max}$.

THEOREM 3.16. *[Uniqueness for large N.] Suppose $\varepsilon < r_{\max}$. Then there exists $\bar{N}(\varepsilon)$ such that there is a* unique *equilibrium threshold for all $N > \bar{N}(\varepsilon)$.*

PROOF. We want to prove that there exists N^* such that for all $N > N^*$, there is no solution to $p(r) = r$ in the interval $(r_b^*(N), r_{\max})$, where recall that $r_b^*(N)$ is the best equilibrium, or smallest solution to this equation in $[r_{\min}, r_{\max})$.

By assumption, $\varepsilon < r_{\max}$.

Now, $p(r)$ is increasing in r for any N (Proposition 3.6), and $p(r_{\max}) = \varepsilon$. So any solution to $p(r) = r$ for $r \in [r_b^*, r_{\max})$ must lie in the subinterval $[r_b^*, \varepsilon)$, since $p(r) < p(r_{\max}) = \varepsilon$ for all $r \in [r_{\min}, r_{\max})$.

Since $\varepsilon < r_{\max}$, $1 - F(\varepsilon) > 0$ by assumption that F is strictly increasing. Let N_1 be the smallest integer such that

$$N_1 > \frac{1}{1 - F(\varepsilon)}.$$

Then, for all $N \geq N_1$,

$$p(\varepsilon, N) = \frac{\varepsilon}{N} \frac{1 - F^N(\varepsilon)}{1 - F(\varepsilon)} < \frac{\varepsilon}{N} \frac{1}{1 - F(\varepsilon)} < \varepsilon,$$

since $\frac{1}{1-F(\varepsilon)} < N_1 \leq N$ by choice of N_1. Therefore, $p(\varepsilon, N) < \varepsilon$ for all $N \geq N_1$.

Next we will choose N large enough so that $p(r, N) - r$ is decreasing in r on $[r_{\min}, \varepsilon]$. For $p(r) - r$ to be decreasing on this interval, we need the derivative of $p(r) - r$ to be negative for all $r \in [r_{\min}, \varepsilon]$, *i.e.*,

$$\frac{\partial}{\partial r}(p(r, N) - r) = \frac{\varepsilon f(r)}{1 - F(r)} \left[\frac{1 - F^N(r)}{N(1 - F(r))} - F^{N-1}(r) \right] - 1 < 0.$$

Since $f(r)$ and $1 - F(r)$ are nonnegative, this is satisfied if

$$\frac{\varepsilon f(r)}{1 - F(r)} \left[\frac{1 - F^N(r)}{N(1 - F(r))} \right] < 1$$

on the interval $[r_{\min}, \varepsilon]$. By assumption $f(r)$ is bounded everywhere on its support, say, by c. Note that the function $\frac{1-F^N(r)}{(1-F(r)^2)}$ is decreasing in r for any fixed N using Proposition 3.6 and the fact that $\frac{1}{1-F(r)}$ is increasing in r. So let N_2 be such that

$$N_2 > \frac{\varepsilon c}{1 - F(\varepsilon)} \left[\frac{1 - F^N(\varepsilon)}{(1 - F(\varepsilon))} \right],$$

so that for all $N \geq N_2$ and $r \leq \varepsilon$, $p'(r, N) < 1$. (Note that we are *not* claiming that the derivative $p'(r)$ is monotone in r—we are simply saying that it is bounded away from 1 for all $r \in [r_{\min}, \varepsilon]$, which again is possible because $F(\varepsilon)$ is bounded away from 1 since $\varepsilon < r_{\max}$ by assumption.)

Now for $N \geq N_2$, we have that the derivative of $p(r) - r$ is negative on $[r_{\min}, \varepsilon]$. Note that $p(r, N) = r$ at $r = r_b^*(N)$ so that $p(r, N) - r = 0$ at $r = r_b^*(N)$, and $p(\varepsilon) - \varepsilon < 0$. Therefore $p(r, N) - r < 0$ for all $r \in (r_b^*, \varepsilon]$ when $N \geq N_2$. But if we choose $N \geq \max(N_1, N_2)$, then this means that there is no equilibrium threshold in $[r_b^*, r_{\max})$, since for $N \geq N_1$ any solution to $p(r, N) = r$ in $(r_b^*, r_{\max})$ must lie in (r_b^*, ε) and for $N \geq N_2$ we showed that there is no such solution.

This proves the result with $\bar{N} = \max(N_1, N_2)$. □

We note, however, the function $\bar{N}(\varepsilon)$ increases with ε, and diverges as $\varepsilon \to r_{\max}$.

3.6. Privacy violation

The equilibrium thresholds r^* we analyze arise from agents who make their equilibrium participation decisions based on the *expected* privacy they will receive, given that all other agents also decide whether or not to participate based on the threshold strategy r^*. (Recall here from §2.2

that comparing *expected* privacy against the requirement does indeed correspond to expected utility maximization in a standard model of privacy.) However, it is possible, in an "unlucky" draw of privacy requirements r_i from the distribution F, that the number of actual participants is smaller than the expected number of participants $\mathbb{E}[\frac{\varepsilon}{n(r^*)}]$, and thus for the agents who do participate to receive a weaker privacy guarantee than they had expected, and possibly even a worse guarantee than their privacy requirement r_i.

While all agents receive privacy that at least meets their requirement r_i *in expectation*, what is the probability that the actual privacy received in a particular instance is worse than an agent's requirement? We say that an individual i with requirement r_i experiences a *privacy violation* if the total number of participants, including herself, is smaller than $\frac{\varepsilon}{r_i}$. In an equilibrium with threshold r^*, the probability that i receives privacy below her requirement is therefore the probability that at most $\frac{\varepsilon}{r_i} - 2$ among the remaining $N-1$ individuals participate, *i.e.*, have privacy requirements $r > r^*$. For the marginal individual with requirement $r_i = r^*$, this probability of a violation (which is the same as the probability that any other participating agent receives privacy below the expected guarantee at equilibrium) is then

$$p_v(F, N, r^*) = \sum_{i=0}^{\frac{\varepsilon}{r_i}-2} \binom{N-1}{i} (1 - F(r^*))^i (F(r^*))^{N-1-i},$$

as this is the probability that exactly 0 of the $N-1$ others participate, plus the probability that exactly 1 participates, all the way up to $\varepsilon/r^* - 2$ others participating, which would leave the marginal individual just shy of her privacy requirement.

If one is considering an equilibrium where $\lim_{N\to\infty} r^* = m > 0$, *i.e.*, the equilibrium threshold is bounded away from zero from below by some constant m (for example, when $r_{\min} > 0$), and where in the limit $F(r^*)$ is bounded away from 1 from above (*i.e.*, at least some individuals participate in the limit) we can upper bound this by

$$p_v(F, N, r^*) \leq \sum_{i=0}^{\frac{\varepsilon}{r_i}-2} \binom{N-1}{i} (F(r^*))^{N-1-i},$$

whose limit behavior is then

$$O\left((N-1)^{\frac{\varepsilon}{m}-2} F(r^*)^{N+1-\frac{\varepsilon}{m}}\right),$$

which goes to 0 as $N \to \infty$, so that the limiting probability of a privacy violation goes to zero as well.

If in the limit r^* is not bounded away from zero, however, the probability of a privacy violation depends much more delicately on the rate of change of r^* and $F(r^*)$ in N. To illustrate this, consider f is the uniform distribution on $[0, 1]$ and $\varepsilon = 0.2$. Then the probability of a privacy violation can be numerically observed to increase up to converge to a non-zero value, approximately 1.75 percent. We discuss the question of providing privacy guarantees that meet the privacy requirements in *every instance*, rather than in expectation, further in §4.2.

4. FURTHER DIRECTIONS

In this paper, we analyzed equilibrium existence and behavior in the game that arises naturally when an analyst proposes a specific computation, and privacy-sensitive agents make simple binary decisions about whether or not to participate based on the proposed computation. Our analysis shows that there can be unexpected behavior in terms of non-monotonicities and multiplicity of equilibria, leading to a first set of qualitative insights—for example, while one would expect that offering to add more noise to the outcome should elicit higher participation, this need not be the case; the results on noise values which support unique equilibria and the large population results

suggest guidelines for an analyst on choosing noise values where these effects do not arise or are mitigated.

A number of immediate questions arise in this setting, the most interesting of which are regarding *design*—how might the game be modified so that outcomes are improved? We discuss some natural directions for further work below.

4.1. Analyst utility

We have so far analyzed the equilibrium expected participation in the game that arises when individuals respond to to an announced computation ε based on their expectations about overall participation and therefore the privacy they will receive. Consider now an analyst who wants to choose ε to optimize his utility, to trade off the two sources of inaccuracy in her noisy computation: sample error due to only receiving data from a subset of the population (recall that we assume individuals' privacy requirements and their data are uncorrelated, so there is no concern of non-response bias), and the direct inaccuracy introduced by the noise he adds. The equilibrium analysis we carried out, particularly regarding existence and uniqueness of equilibria, are likely salient to an analyst attempting to optimize his utility, since they provide a basis for estimating the nature and extent of participation as a function of the announced noise ε.

To actually formulate an optimization in terms of ε, however, the analyst needs to choose an appropriate metric to optimize, which will depend on his specific computation and its desired application; in addition, there may also be setting-specific bounds on the noise level. He must also specify how much disutility he incurs in the (typically non-zero-probability) event that no individuals participate, so that his computation simply cannot be run. Also, for regimes of (F, ε, N) where multiple equilibria exist, there are different choices (worst-case or best-case, to name just two extremes) as to how he should analyze his expected utility; or, perhaps the analyst wishes to restrict his optimization to ranges of ε where there exists a unique equilibrium. Our equilibrium analysis provides the first inputs necessary for formulating such an optimization; however, a full formulation requires a number of modeling choices which are application-dependent, and is left to future work.

4.2. Sharing information

A natural, and simple, tool available to an analyst who acts also as a mechanism designer, is to share *information* about agents' participation with other agents. As we discussed in §3.6, agents in our model only receive privacy that meets their privacy requirements *in expectation*, and the probability that the privacy received in a particular instance meets an agent's requirement is less than one for any finite population size. An analyst who wishes to ensure that all individuals receive privacy guarantees that meet their requirements instance-by-instance, and not just in expectation, might consider sharing information about other agents' participation, for instance by keeping a constantly updated public posting of the number of individuals who have elected to participate so far. Individuals may then decide whether or not to participate based on whether the number of current participants is adequate to give them the privacy guarantees they desire. Here, one can imagine two natural situations, one where individuals arrive sequentially and have a single opportunity to participate (given the information about the current number of participants), or where they are 'always present' (such as, for example, on a website or application they use regularly) and may choose to participate at any point. Note that if agents choose to use no distributional information and make their participation decisions purely based on the number of current participants, assuming there will be no further participants beyond themselves, each agent will definitely receive privacy that meets, and is likely better than, her privacy requirement. But what is the effect of such publicly posted information on equilibrium participation?

In either model of agent arrival, note that if $\varepsilon \geq r_{\max}$, there is no participation at all, since no agent will be willing to be the first to participate. However, for $\varepsilon < r_{\max}$ (so that $1 - F(\varepsilon) > 0$), expected participation in both models is very good: we can show that the expected number of participants diverges in the limit as N goes to infinity, and grows as $\Theta(N)$, since the probability that each individual participates is at least a constant. The question of how exactly participation grows,

and how participation with sharing such information compares against participation in the Bayes-Nash setting—both in the limit and for finite N—is an interesting question, and addresses the role of information as a design choice in incentivizing participation from a privacy sensitive population.

4.3. Eliciting higher participation

A natural question to ask is whether there are other, more sophisticated mechanisms for eliciting participation that lead to a larger subset of the population participating in equilibrium than that achieved by simply announcing the computation and letting agents make choices based on their beliefs about other agents' privacy requirements. Sharing information, as we discussed above, is one such scheme. But other schemes are possible as well, even without attempting to explicitly elicit agents' privacy requirements[6]. For instance, a mechanism might offer different levels of added noise to different subsets of the population, or progressively announce different levels of added noise as agents gradually announce participation. There is a vast space of mechanisms possible here, including (rather than announcing computations) announcing privacy guarantees that improve progressively as more and more agents sign up for the computation (so that the analyst is not promising to add an unacceptably high level of noise at any stage in the mechanism). We note also that another possible method for improving outcomes is equilibrium selection mechanisms for settings with multiple equilibria, to induce the most desirable equilibrium with highest expected participation.

In this vein, it is also natural to ask how participation, both in our simplest setting of merely announcing the computation, as well as in other more sophisticated mechanisms, compares with natural benchmarks. For example, one natural benchmark to consider is the best possible participation rate (or analyst utility) that could be achieved by a mechanism constrained to offer all individuals the same privacy guarantee, if individuals' privacy requirements were common knowledge—this corresponds, for any given N, do finding the largest k such that the expected number of individuals with requirements at or above ε/k is at least k. A stronger benchmark is the value of the metric under the best omniscient mechanism that performs some arbitrary aggregation of the individuals' data (such as a weighted sum rather than a mean) before adding noise, in a manner that respects the privacy requirements of all individuals but does not give them all the same privacy guarantee.

4.4. Other directions

There are also a number of other questions both in the space of design for achieving better outcomes, and otherwise. In the previous two subsections, we took the distribution F of thresholds as given— we assumed that agents' privacy requirements could not be modified, for instance, by offering them additional value (beyond what they already perceive) from participating in the computation. Suppose, however, that there are some agents who could be incentivized to participate, either with a monetary reward or some other form of benefit such as superior personalization. Can such agents be used, especially if participation levels are shared with other agents, to set off a cascade of participation? A natural extension of this problem is to a network setting, where agents' data is only used in conjunction with data from their neighborhood, such as in mobile health applications leveraging social network data. The large literature on adoptions and cascades in network settings, too large to properly cite here, including recent work on strategies for maximizing adoption in social networks, is technically related to this question and could potentially be used to understand how to seed to achieve high final participation, for instance in applications involving data sharing on social networks.

Finally, we assume that there is no correlation between an agent's privacy requirement and her private data— while this is not unjustifiable in several online settings, there are also several other settings such as computations on medical data where this assumption is unlikely to hold. Explor-

[6] As discussed earlier, attempting to elicit agents privacy requirements might not always be practical because a typical person might be unable to state a quantitative privacy requirement if asked, but is likely to be able to say whether or not she is willing to share her data in any given scenario.

ing settings where privacy requirements are correlated with thresholds remains a fascinating but challenging direction for further work.

REFERENCES

APERJIS, C. AND HUBERMAN, B. 2012. A market for unbiased private data: Paying individuals according to their privacy attitudes. *Available at SSRN 2046861*.

BAGNOLI, M. AND LIPMAN, B. 1989. Provision of public goods: Fully implementing the core through private contributions. *The Review of Economic Studies 56,* 4, 583–601.

BRANDT, F. AND SANDHOLM, T. 2008. On the existence of unconditionally privacy-preserving auction protocols. *ACM Transactions on Information and System Security (TISSEC) 11,* 2, 6.

CHEN, Y., CHONG, S., KASH, I., MORAN, T., AND VADHAN, S. 2011. Truthful mechanisms for agents that value privacy. *arXiv preprint arXiv:1111.5472*.

DANDEKAR, P., FAWAZ, N., AND IOANNIDIS, S. 2012. Privacy auctions for recommender systems. In *Workshop on Internet Economics (WINE)*.

DINUR, I. AND NISSIM, K. 2003. Revealing information while preserving privacy. In *Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*. ACM Press New York, NY, USA, 202–210.

DWORK, C., MCSHERRY, F., NISSIM, K., AND SMITH, A. 2006. Calibrating noise to sensitivity in private data analysis. In *Proc. Theory of Cryptography Conference*. 265–284.

DWORK, C. AND NISSIM, K. 2004. Privacy-preserving datamining on vertically partitioned databases. In *In CRYPTO*. Springer, 528–544.

EASLEY, D. AND KLEINBERG, J. 2010. *Networks, Crowds, and Markets*. Cambridge University Press.

EDLIN, A., FARRELL, J., AND SEGAL, I. 1998. The Edlin, Farrell, Segal mechanism for the establishment of the Berkeley Electronic Press. Personal communcation from Preston McAfee.

FEIGENBAUM, J., JAGGARD, A., AND SCHAPIRA, M. 2010. Approximate privacy: foundations and quantification. In *Proceedings of the 11th ACM conference on Electronic commerce*. ACM, 167–178.

FEIGENBAUM, J., MITZENMACHER, M., AND ZERVAS, G. 2012. An economic analysis of user-privacy options in ad-supported services. *arXiv preprint arXiv:1208.0383*.

FLEISCHER, L. AND LYU, Y. 2012. Approximately optimal auctions for selling privacy when costs are correlated with data. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 568–585.

FUDENBERG, D. AND TIROLE, J. 1991. *Game Theory*. MIT Press.

GHOSH, A. AND ROTH, A. 2011. Selling privacy at auction. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM, 199–208.

GKATZELIS, V., APERJIS, C., AND HUBERMAN, B. 2012. Pricing private data. *Available at SSRN 2146966*.

HUANG, Z. AND KANNAN, S. 2012. The exponential mechanism for social welfare: Private, truthful, and nearly optimal.

ISAACMAN, S., IOANNIDIS, S., CHAINTREAU, A., AND MARTONOSI, M. 2011. Distributed rating prediction in user generated content streams. In *Proceedings of the ACM RecSys*.

KEARNS, M., PAI, M., ROTH, A., AND ULLMAN, J. 2012. Mechanism design in large games: Incentives and privacy. *arXiv preprint arXiv:1207.4084*.

KLEINBERG, J., PAPADIMITRIOU, C., AND RAGHAVAN, P. 2001. On the value of private information. In *Proceedings of the 8th conference on Theoretical aspects of rationality and knowledge*. Morgan Kaufmann Publishers Inc., 249–257.

LAUDON, K. 1993. Markets and privacy. *Information Systems Working Papers Series, Vol.*

LEYTON-BROWN, K. AND SHOHAM, Y. 2008. *Essentials of Game Theory: A Consise, Multidisciplinary Introduction*. Morgan and Claypool.

LI, C., LI, D., MIKLAU, G., AND SUCIU, D. 2012. A theory of pricing private data. In *Interna-*

tional Conference on Database Theory (ICDT).

LIGETT, K. AND ROTH, A. 2012. Take it or leave it: Running a survey when privacy comes at a cost. In *Workshop on Internet Economics (WINE)*.

MCSHERRY, F. AND TALWAR, K. 2007. Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*. IEEE, 94–103.

NAOR, M., PINKAS, B., AND SUMNER, R. 1999. Privacy preserving auctions and mechanism design. In *Proceedings of the 1st ACM conference on Electronic commerce*. ACM, 129–139.

NISSIM, K., ORLANDI, C., AND SMORODINSKY, R. 2012a. Privacy-aware mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 774–789.

NISSIM, K., SMORODINSKY, R., AND TENNENHOLTZ, M. 2012b. Approximately optimal mechanism design via differential privacy. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*. ACM, 203–213.

ROTH, A. AND SCHOENEBECK, G. 2012. Conducting truthful surveys, cheaply. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 826–843.

SAMUELSON, P. 1983. *Foundations of Economic Analysis*. Harvard University Press.

VARIAN, H. 1994. Sequential contributions to public goods. *Journal of Public Economics 53,* 2, 165–186.

XIAO, D. 2013. Is privacy compatible with truthfulness? In *Innovations in Theoretical Computer Science (ITCS)*.