

Information Design for Differential Privacy

Ian M. Schmutte and Nathan Yoder*

June 18, 2025

Abstract

Firms and statistical agencies must protect the privacy of the individuals whose data they collect, analyze, and publish. Increasingly, these organizations do so by using publication mechanisms that satisfy *differential privacy*. We consider the problem of choosing such a mechanism so as to maximize the value of its output to end users. We show that mechanisms which add conditionally independent noise to the statistic of interest — like most of those used in practice — are never without loss when the statistic is a sum or average of magnitude data (e.g., income). However, we also show that adding conditionally independent noise is *always* optimal when the statistic is a count of data entries with a certain characteristic, and the underlying database is drawn from a symmetric distribution (e.g., if individuals' data are i.i.d.). When, in addition, data users have supermodular payoffs, we show that the simple *geometric mechanism* is always optimal by using a novel comparative static that ranks information structures according to their usefulness in supermodular decision problems.

Keywords: Bayesian persuasion, information acquisition, comparison of experiments

JEL Codes: D83, D81, C81

*Schmutte: University of Georgia, Terry College of Business, John Munro Godfrey Sr. Department of Economics; E-mail: schmutte@uga.edu. Yoder: University of Georgia, Terry College of Business, John Munro Godfrey Sr. Department of Economics; E-mail: nathan.yoder@uga.edu. An extended abstract of this paper appeared in *EC '22: Proceedings of the 23rd ACM Conference on Economics and Computation*. The authors wish to thank Gary Benedetto, Mark Fleischer, Kevin He, R. Vijay Krishna, Meg Meyer, Marco Ottaviani, John Quah, Marzena Rostek, Alex Smolin, Peter Norman Sørensen, Nikhil Vellodi, Juba Ziani, seminar participants at Bank of Canada, Concordia University, Emory, Florida State, Paris School of Economics, Penn State, and Yale, and conference attendees at NASMES 2022, EC'22, the 2022 ASU Theory Conference, SAET 2023, NBER Data Privacy 2024, and ESSET 2024 for their helpful comments, as well as Cole Wittbrodt for excellent research assistance. Figures created with Wolfram Mathematica. Both of the authors acknowledge summer support from the University of Georgia through a Terry-Sanford Research Award. Schmutte is grateful for financial support from the Bonbright Center for the Study of Regulation.

1 Introduction

To help people make better decisions, statistical agencies collect and then disseminate information about data. For instance, information about the prevalence of disease can affect decisions that impact public health, while information about the level of unemployment can affect the decisions of participants in the labor market. But disseminating information about disease prevalence or unemployment also necessarily reveals information about whether any one person is sick or unemployed (Dinur and Nissim, 2003). In order to address public concerns about data privacy, and the emerging policy and legal responses to those concerns, statistical agencies and other data providers must find ways to release information that is as informative as possible about important population characteristics, while also protecting the privacy of individuals.

This has prompted technology firms like Google, LinkedIn, and Uber, and statistical agencies like the U.S. Census Bureau, to adopt mechanisms for publishing data that provide formal, mathematically provable guarantees of privacy.¹ Frequently, these guarantees take the form of *differential privacy* (Dwork et al., 2014; Fioretto et al., 2024). This privacy criterion — which we describe in detail — requires a data provider to randomize its output so as to provide an explicit quantitative bound on the amount that a change in any individual entry in the data can change the probability of producing any given output.

Many different *publication mechanisms* — or (stochastic) maps from true data to published output — have been proposed for publishing statistics in a way that satisfies differential privacy (e.g., Ghosh et al. (2012); Geng and Viswanath (2015)). Because differential privacy limits the amount of information that can be revealed about individuals, these generally aim to disclose information about aggregate quantities — i.e., population statistics — such as the proportion of respondents in the data that are unemployed, or the average of those respondents’ incomes. Comparisons between publication mechanisms have focused on their *accuracy* — the expected value of (a function of) the distance between the published output and the true value of the statistic.

However, there is much less clarity about how to maximize the *value to data users* of published statistics while preserving privacy. Our paper sheds light on this question by introducing an information design approach to differential privacy. In particular, we endow data users with payoffs that depend on their actions and the true value of a population statistic, and explicitly model the way that their decisions about the former depend on beliefs about the latter. We assume that data users are Bayesian, and update these beliefs after observing the output of a signal about the contents of the database from which the statistic is drawn — i.e., a publication mechanism. Then, we consider the problem of a data provider who chooses that signal under commitment — i.e., before accessing the data — so as to maximize the data users’ welfare, subject to a differential privacy constraint.

This is essentially the problem faced by many statistical agencies and technology firms in the real world. The U.S. Census Bureau has announced its intention to adopt formal privacy requirements, including differential privacy, as its primary approach to disclosure limitation.² Apple uses differ-

¹See Guevara (2019) (Google), Rogers et al. (2020) (LinkedIn), Near (2018) (Uber).

²Abowd et al. (2020) describes how formal privacy is to be applied in several flagship products, including the Decennial Census, American Community Survey and the Economic Census. For many products, the exact privacy requirements

ential privacy to protect individuals’ privacy in Apple’s own analysis of aggregate user behavior (Apple Inc., 2017). During the early COVID-19 pandemic, when Facebook published information on aggregate mobility patterns to aid public health researchers and guide policy, they protected individuals with a differential privacy guarantee (Dow et al., 2020). Crucially, offering a meaningful differential privacy guarantee requires data providers like these to commit to a mechanism without reference to the underlying data.³

The data provider in our model faces a special case of a more general problem: choosing a *signal* under *commitment* to influence the *decision* of a Bayesian agent who observes it. This general problem is the focus of the extensive literature on information design (e.g., Kamenica and Gentzkow (2011); Taneva (2019); Kamenica (2019)). By introducing key tools from this literature to the study of differential privacy — and developing new tools that apply to information design more broadly — we provide new insights that have direct, practical applications. We describe these — our main results — in the context of an applied example.

Example 1 (School Planning). A school district needs to choose how many slots a to create in a universal pre-K program for the next three years. It would like to choose a number that is as close as possible to the actual number ω of children living in the district who are under 4, and hence will become eligible to enroll during that time: Empty slots will not be paid for by the state, and must instead come out of the district’s budget, whereas if there are not enough slots, the district will be penalized by the state. For simplicity — and since it facilitates a more direct comparison with the existing literature on accuracy in differential privacy — we assume that its loss function is simply squared error: $(\omega - a)^2$.

Since the children that will enroll are not currently in school, the district must rely on statistics published by the U.S. Census Bureau to make its decision. But the Census Bureau has statutory obligations to preserve the privacy of the respondents (here, households) in its database. Suppose that it has chosen to publish information about ω in a way that is differentially private, with a “privacy loss budget” (the bound on the amount that a change in a single data record can change the log probability of producing an output) of $\epsilon = 1$. The records θ_i in its database each give the number of children under 4 in a household i , which varies between 0 and 2. Thus, a change in a data record can change ω by at most $\Delta = 2$.

In keeping with actual practice, we assume that the Bureau chooses to satisfy differential privacy by *adding conditionally independent noise* to the true value of ω ; i.e., publishing $s = \omega + \eta$, where η is a random variable whose distribution is independent of the structure of the database, conditional

are currently under discussion, but Census has generally used differential privacy and closely-related privacy concepts as a starting point (Vilhuber and Schmutte, 2016). Among those products for which privacy guarantees have been finalized, the Census Bureau’s Post-Secondary Employment Outcomes data provides the same differential privacy guarantee we consider here, while the 2020 Decennial Redistricting File offers a guarantee of *zero-concentrated differential privacy*.

³For example, the U.S. Census Bureau designed the publication mechanism for the 2020 Redistricting Files using demonstration data based on prior census data; US Census Bureau (2021) gives a timeline summarizing this process. See Section 2 for a detailed discussion of why using the data to develop or calibrate the mechanism can make the mechanism’s privacy guarantee meaningless.

on ω .⁴ In particular, suppose that as with several of its other data products,⁵ it chooses to add noise that has a *two-sided geometric* (or equivalently, discretized Laplace) distribution with parameter $\epsilon/\Delta = 1/2$.⁶ Its choice to do so is motivated by the results of Ghosh et al. (2012), who show that if data users can postprocess the published output (i.e., remap some values of s to others), and the population statistic ω is a count of records with some characteristic, then among all ways to add conditionally independent noise, this one minimizes the expected squared error $(s - \omega)^2$ of the published output relative to the true statistic.^{7,8} Ghosh et al. (2012) call this the *geometric* publication mechanism.

We evaluate the performance of this mechanism — in terms of the value of its output to the data user — in a geography with $N = 40$ households with i.i.d. data records. In particular, we form our prior about the number of children under 4 in a household using the distribution observed in the 2019 5-Year American Community Survey within Athens-Clarke County, Georgia:

$$P(\theta_i = 0) = 0.89, \quad P(\theta_i = 1) = 0.09, \quad P(\theta_i = 2) = 0.02.$$

In this context, when the school district chooses a after seeing the output of the geometric mechanism, its expected squared error is 3.37. But despite the optimality result of Ghosh et al. (2012), we can do better. In fact, Proposition 1 shows how to design a mechanism that lets the district achieve an expected loss of 2.48 — about 26% better. This represents an *inward shift* of the privacy-accuracy frontier described in Abowd and Schmutte (2019): In addition to improving the accuracy of published output while holding privacy loss constant, we can also use Proposition 1 to decrease privacy loss ϵ while holding accuracy (in the form of expected squared error) constant.

The reason this is possible is that, as we show in Theorem 1, *in this setting, adding conditionally independent noise is generally not the optimal way to satisfy the differential privacy constraint*.⁹ Instead, we can do better by allowing the distribution of published output to depend on the entire database, rather than just the true population statistic ω . This fact is somewhat counterintuitive, since the data users' payoff depends only on the database through ω . Indeed, we are not aware of any discussion of this point in the literature on differential privacy. As we show, taking an information

⁴In practice, η usually depends on ω only through truncation of the possible values of s ; e.g., the truncated ϵ -geometric mechanism described in footnote 33.

⁵For instance, the U.S. Census Bureau uses the geometric mechanism to publish industry-specific counts of employed graduates of postsecondary education programs (Foote et al., 2021). It is also used in the private sector: in the early COVID-19 pandemic, Facebook used its continuous analogue, the Laplace mechanism, to publish counts of individuals that did not leave a small area over the course of the day (Dow et al., 2020).

⁶Formally, for each ω , the distribution of the noise term is given by

$$P(\eta = x) = \frac{1 - e^{-\epsilon/\Delta}}{1 + e^{-\epsilon/\Delta}} e^{-x\epsilon/\Delta} = \frac{1 - e^{-1/2}}{1 + e^{-1/2}} e^{-x/2}.$$

⁷In fact, it minimizes the expected value of any decreasing function $\ell(\omega, |s - \omega|)$ of the error.

⁸In particular, if they are both differentially private for the same privacy loss budget ϵ , the geometric mechanism outperforms the truncated *Gaussian mechanism* — in which e follows a normal distribution, and the output s is truncated to lie in the range of possible values of ω — used to publish data from the 2020 Decennial Census.

⁹In settings with magnitude data, we can also potentially do better by using a mechanism that adds conditionally independent noise that is not geometric; see, e.g., Geng and Viswanath (2015). But at least in this example, the gains are much smaller: The best we can do is reduce expected squared error to 3.27, an improvement of only 3% over the geometric mechanism.

design approach to the problem — and thus thinking of differential privacy as a constraint on the *distributions of posterior beliefs* about ω that can be induced by the data publisher — allows us to be precise about the sense in which adding conditionally independent noise is not without loss.

But, as we show, there are other settings in which adding conditionally independent noise *is* without loss. Moreover, we show that the geometric mechanism is the optimal way to do so in a very broad class of settings in which data users’ payoffs need not bear any resemblance to expected squared error.¹⁰

To illustrate this, suppose that instead of how many places to create in its program, the district must determine how many school buses to purchase for it. Instead of the number of children who are likely to enroll in the program, the optimal number of bus routes depends on the *number of households* with such children. That is, the population statistic ω is a *count* of *binary* data entries, whereas in the number-of-places problem it was the *sum* of *magnitude* data. As we show in Theorem 2, as long as it views households as anonymous — i.e., no household is more likely to have children than any other, though their statuses may be correlated — this makes it optimal to ensure differential privacy by adding conditionally independent noise.

But how should it add such noise? Suppose that the district seeks to minimize the sum of the average number of households per bus route and the district’s expenditure on buses: when there are ω households with children under 4 and the district purchases a buses, its payoff is $u(a, \omega) = -\frac{\omega}{a} - ca$, for some $c > 0$.

In this decision problem, the district does not seek to minimize squared error, or any of the other loss functions considered in Ghosh et al. (2012). But even though it does not want its action to *match* the statistic ω , a rightward shift in its belief about ω still causes its optimal action to increase; i.e., its payoff is *supermodular*. In Theorem 3, we prove that this is the key feature that makes the geometric mechanism optimal.

The kind of decision problems considered in Example 1 — those in which a data user’s decision depends on their belief about a population statistic — are common in practice. For instance, if an individual believes there is a greater incidence in her community of a pandemic pathogen, she may be more likely to take precautions such as avoiding indoor dining or wearing an N95 mask. If a restaurateur believes that the average income in a neighborhood is higher, he may decide to implement a more upscale menu. And if a municipal government believes that a greater proportion of its constituents are unemployed, it may increase the tax incentives that it offers potential employers. Moreover, as in Example 1 (and each of the applications listed above), data users frequently face problems that are *supermodular*; i.e., that feature complementarity between higher actions and higher beliefs about the population statistic.¹¹

Our results offer guidance to a data provider who views these problems as representative use cases for the information it publishes. If the population statistic of interest is a sum of magnitude data, adding conditionally independent noise is generally not optimal (Theorem 1); the data

¹⁰Or, more generally, the kind of loss functions considered in Ghosh et al. (2012) that depend on the (remapped) signal, or action, through its distance from the statistic.

¹¹By higher beliefs, we mean beliefs that are shifted to the right; i.e., higher in the sense of first-order stochastic dominance.

provider can do better with a publication mechanism that depends not only on the *average* of the data, but on the sample *distribution* more generally. On the other hand, if it is a count of binary data from respondents that are ex ante anonymous, adding conditionally independent noise is without loss (Theorem 2) — and in the supermodular case, the geometric mechanism is always optimal (Theorem 3).

We also provide a general comparative static on the comparison of information structures that we use to establish Theorem 3, but which is of independent interest. In particular, we define a dominance criterion on information structures about ordered states of the world, and show in Theorem 4 that it ranks signals higher when they are more useful to decision makers with supermodular problems. This notion of dominance is novel: Because of the nature of the data provider’s problem, Theorem 3 cannot appeal to existing tools for ranking experiments according to their usefulness to a class of decision makers.¹² Instead, we introduce *Uniform-Peaked Relative Risk (UPRR) dominance*, which ranks information structures higher when the posteriors they induce each concentrate relatively more mass around a *peak* (e.g., their mode). Intuitively, we might expect UPRR-dominant information structures to be more desirable in settings where the marginal benefit of taking a higher action is increasing in the state: There, the opportunity cost of choosing an action when it is sub-optimal increases as the state gets further away from the region where that action is optimal, so if posteriors are more concentrated, the costs associated with such mistakes should be lower ex ante. Theorem 4 confirms this intuition, showing that information structures higher in the UPRR order are more useful for maximizing supermodular payoffs. Since the geometric mechanism is UPRR-dominant over *all* differentially private data publication mechanisms that depend only on the true value of the population statistic (Lemma 3), Theorem 3 follows.

Related Literature Our paper complements a growing body of research addressing the challenge posed by Abowd et al. (2019) for economists to develop analytical tools that prioritize information quality in privacy-preserving data publication. Several recent papers treat the publication mechanism as fixed, and ask how the data provider should set the privacy budget, both as a theoretical matter (Abowd and Schmutte, 2019; Hsu et al., 2014; Echenique and He, 2021) and as a practical matter (Chetty and Friedman, 2019). By contrast, we treat the privacy requirement as fixed, and try to find an optimal publication mechanism. Other authors focus on how property rights in data should be assigned (Jones and Tonetti, 2020; Arrieta-Ibarra et al., 2018), or the design of privacy-preserving mechanisms (Pai and Roth, 2013; Eilat et al., 2021).

Our paper also relates to prior work in computer science that examines the performance of differentially private publication mechanisms (e.g., Geng and Viswanath (2015); Koufogiannis et al. (2015)). This literature generally focuses on maximizing *accuracy* (i.e., minimizing some nondecreasing function of the distance between the published output and the true population statistic) and *oblivious* mechanisms — those that add conditionally independent noise to the population statistic.

¹²In particular, the information structures satisfying Corollary 3’s characterization of the solution to the data provider’s problem — i.e., those which induce posteriors that are extreme points of the constraint set — are not Blackwell-comparable (making Blackwell’s theorem unavailable) and do not necessarily induce an MLRP-ordered set of posterior beliefs (making comparative statics on Lehmann (1988) accuracy unavailable).

Among these papers, two are closest to our work. First, as discussed in Example 1, Ghosh et al. (2012) show that when the statistic of interest is the count of binary data, and the data user can use his prior to “remap” the mechanism’s output, the geometric mechanism maximizes accuracy among all oblivious mechanisms. When, in addition, data users have priors about the population statistic but view the database as ambiguous, they show the geometric mechanism also outperforms mechanisms that are not oblivious. In contrast, our Theorem 3 shows that with binary data, the geometric mechanism is the optimal differentially private oblivious mechanism *whenever data users have supermodular decision problems*, and that using an oblivious mechanism is always optimal with binary data and a symmetric prior (Theorem 2), but is generally not with magnitude data (Theorem 1). Second, Brenner and Nissim (2014) show that when the statistic of interest is a sum of magnitude data, there is no *single* oblivious mechanism that maximizes accuracy for *every* nondecreasing loss function and *every* prior, the way that Ghosh et al. (2012) show the geometric mechanism does in the binary data case. In contrast, our Theorem 1 compares the *entire class* of oblivious mechanisms to those that depend on the database more generally, and shows that the latter can outperform the former with magnitude data.

This paper follows other recent work that uses the tools of information design to study privacy.¹³ Many of these papers focus on questions about the optimal *amount* of privacy protection, generally using concepts other than differential privacy: Ichihashi (2020) considers the way that consumers will choose to flexibly disclose information to sellers; Ali et al. (2022) and Hidir and Vellodi (2021) consider the optimal regulation of such disclosure. We focus instead on a question similar to those considered by He et al. (2023) and Strack and Yang (2024): how should information be published, given the presence of a privacy constraint? Our approach differs from those papers, however, because differential privacy only protects sensitive information *imperfectly*. All of these (including our paper) are part of a larger body of research on the economics of privacy; see Acquisti et al. (2016) for a survey.

Finally, our Theorem 4 contributes to the literature on the comparison of experiments following Blackwell (1953), whose ordering ranks signals according to their usefulness for *any* decision maker. Lehmann’s (1988) *accuracy* ordering ranks information structures so that those which are more accurate are more useful to decision makers whose payoffs are single-crossing in actions and states (Persico, 1996). Quah and Strulovici (2009) extend this result to the more general case where payoffs form an *interval dominance order* family. Athey and Levin (2018), on the other hand, consider more general orders on information structures defined by their usefulness in decision problems satisfying various monotonicity conditions, and obtain Lehmann (1988)’s order as a special case. While the supermodular decision problems to which Theorem 4 applies necessarily have the single-crossing property, our result takes a different approach from the literature following Lehmann (1988): Instead of being limited to information structures whose posteriors are ordered by the monotone likelihood ratio property, Theorem 4’s scope extends to *any* pair of signals which are ranked in the UPRR order.¹⁴

¹³Our characterization of the data provider’s problem is also related to other work on information design with constraints on the set of posteriors, e.g., Doval and Skreta (2021); Le Treust and Tomala (2019); Matysková (2019).

¹⁴Recall that one belief is higher than another in the MLRP order — or in the language of Quah and Strulovici (2009),

The paper proceeds as follows. Section 2 formally describes the setting we consider and the meaning of differential privacy within it. Section 3 shows that the data provider faces a constrained information design problem (Proposition 1) and characterizes its solution in the general case (Corollary 1). Section 4 asks when this problem’s dimensionality can be reduced by restricting attention to mechanisms that add conditionally independent noise to the true population statistic, and shows that we cannot do so with magnitude data (Theorem 1) but we can with binary data, so long as respondents are anonymous (Theorem 2). Finally, Section 5 gives our optimality result for supermodular problems with binary data (Theorem 3) and the comparative static on information structures that it relies on (Theorem 4). All proofs are given in Appendix B.

2 Setting

There is a state of the world, or *database*; a data provider, or *designer*, who chooses a mechanism for releasing information about the database; and a *decision maker*, who uses that information to make better decisions.¹⁵

Data A database is a list $\theta = (\theta_1, \dots, \theta_N)$ of N observations $\theta_i \in \{0, 1, \dots, T\}$. Each observation represents the *type* of an individual or household (a *respondent*). We say that the database contains *binary data* when $T = 1$ and *magnitude data* when $T > 1$.

The set of possible databases is thus $\Theta = \{0, 1, \dots, T\}^N$. The decision maker and the designer have a common prior $\pi_0 \in \Delta(\Theta)$ over this set, with full support: $\text{supp}(\pi_0) = \Theta$.^{16,17}

Data Publication and Differential Privacy Before observing the database, the designer chooses a *data publication mechanism* (S, m) , where S is some countable set of outputs and $m : \Theta \rightarrow \Delta(S)$ maps each database to a distribution over outputs in S .¹⁸ The decision maker then observes the realization of $m(\cdot | \theta)$. Without loss, S only includes outputs that the mechanism can actually generate: for each $s \in S$, $m(s | \theta) > 0$ for some $\theta \in \Theta$.

The designer’s chosen mechanism must guarantee the privacy of each respondent’s type by limiting the amount of information that can be revealed about it. In particular, for some *privacy loss* $\epsilon > 0$, the mechanism must be ϵ -*differentially private*.

an *MLR-shift* of the latter belief — if the ratio of the probability placed on higher state to the probability placed on a lower state (i.e., the likelihood ratio of those states) is greater under the first belief than under the second.

¹⁵We assume a single decision maker for simplicity: All of our conclusions are unchanged if there are multiple decision makers and the designer seeks to maximize their total welfare, so long as their decisions do not have externalities on other decision makers, i.e., each DM’s payoff only depends on their own action and the population statistic. (In that case, the problem is the same as if there were a single agent who chooses a vector of actions for the decision makers in order to maximize the sum of their payoffs.)

¹⁶For a set S , we denote its convex hull by $\text{conv}(S)$, its cardinality by $|S|$, its set of extreme points by $\text{ext}(S)$, and the set of Borel probability measures on S by $\Delta(S)$.

¹⁷One motivation for the differential privacy criterion is that it does not depend on assumptions about the background knowledge of agents (“attackers”) that might wish to use the mechanism output to learn about a respondent’s type for malicious purposes. We emphasize that the common prior assumption need not apply to such malefactors, who are not explicitly considered in our model.

¹⁸The assumption that S is countable (or even finite) is without loss in our setting — see Corollary 1(iic).

Definition (ϵ -Differential Privacy (Dwork et al., 2006)). A data publication mechanism (S, m) is ϵ -differentially private if for all outputs $s \in S$ and all databases $\theta, \theta' \in \Theta$ that are adjacent, in the sense that they differ in at most one entry (i.e., such that for some $i \in \{1, \dots, N\}$, $\theta_{-i} = \theta'_{-i}$), we have¹⁹

$$\left| \log \left(\frac{m(s|\theta')}{m(s|\theta)} \right) \right| \leq \epsilon. \quad (1)$$

Differential privacy limits the amount, in log terms, that changing a single respondent's type can change the distribution of the data publication mechanism's realizations.²⁰ This criterion is prior-independent: Whether a mechanism is differentially private does not depend on π_0 . Hence, even if other agents ("attackers") might have more information about the contents of the database than the designer does, a differential privacy guarantee is still valid. In fact, if an agent views types as independent *a priori*, differential privacy is equivalent to a limit on the amount that observing the mechanism's output can cause her to shift her beliefs about an individual respondent's type θ_n , regardless of how much information she already has about the other respondents' types. (See Proposition S.1 in the Online Appendix.)

The Decision Maker After observing the realization of $m(\cdot|\theta)$, the decision maker takes an action a from some compact set A . His Bernoulli payoffs from doing so depend on θ only through the *average* type of the respondents. Hence, since N is fixed, his Bernoulli payoffs can be written as a function $u : A \times \Omega \rightarrow \mathbb{R}$ of his action a and the *population statistic* $\omega_\theta \equiv \sum_{n=1}^N \theta_n$, where $\Omega \equiv \{0, \dots, NT\}$; we assume that u is continuous. Thus, if the decision maker has posterior belief $\pi \in \Delta(\Theta)$ after observing the output of the data publication mechanism, his interim payoffs are given by $v(\pi) \equiv \max_{a \in A} E_\pi u(a, \omega_\theta)$.

The Designer's Objective When choosing the data publication mechanism, the designer acts to maximize the expected value of the decision maker's payoffs. Formally, letting $\pi_s^m \in \Delta(\Theta)$ denote the posterior belief of a decision maker who observes s from the data publication mechanism (S, m) , the designer solves

$$\max_{(S, m)} \left\{ \sum_{\theta \in \Theta} \sum_{s \in S} v(\pi_s^m) m(s|\theta) \pi_0(\theta) \text{ s.t. } \left| \log \left(\frac{m(s|\theta')}{m(s|\theta)} \right) \right| \leq \epsilon \quad \begin{array}{l} \forall s \in S, \forall \theta, \theta' \text{ s.t.} \\ \exists n, \theta_{-n} = \theta'_{-n} \end{array} \right\}. \quad (2)$$

Discussion Our model captures a widely confronted data publication problem in very general terms. The data provider (designer) chooses how the data should be published, and commits to those choices after the structure of the data is determined, but before the actual data are collected.

¹⁹In the differential privacy literature, vectors like these that differ in only one entry are commonly referred to as *neighboring databases*.

²⁰Because there is no uncertainty about the number of respondents in our model, our definition of differential privacy is referred to as *bounded* differential privacy (Kifer and Machanavajjhala, 2011), as distinct from *unbounded* differential privacy, which bounds the change to the signal distribution from adding or removing a respondent from the database. Since mechanisms that satisfy the latter criterion must take as argument databases of multiple sizes, the problem of designing them is somewhat different than the problem we consider. However, the key observation behind Theorems 1 and 2 — with magnitude data (but not binary data), satisfying differential privacy requires oblivious mechanisms to treat databases similarly whenever they have similar population statistics, not just whenever they are adjacent — is still true in the unbounded case, and so we conjecture that those results have counterparts for unbounded DP.

The data provider is also under an obligation to guarantee differential privacy at a fixed level. With binary data, our model corresponds precisely to any case where the data provider will publish information about the number of people in some category — e.g., the number of COVID-positive individuals within a college dormitory, or the number of unemployed workers in a specific occupation within a state. Such statistics are generally called population or frequency counts, or — in the computer science literature — counting queries.²¹ With magnitude data, on the other hand, our model corresponds to cases where the data provider will publish information about the sum or average of some quantitative characteristic — e.g., income, years of schooling, or length of unemployment. These statistics are generally referred to as population totals or averages, or in the computer science literature, sum queries.

The notion of differential privacy that we consider is standard. But unlike most of the literature, we explicitly consider the problem of *choosing* a publication mechanism that satisfies it. This raises the possibility that privacy might be compromised by the designer’s choice of mechanism, rather than by the mechanism’s output. For instance, if the designer were to learn the contents of the database θ before choosing a mechanism, she could perfectly reveal θ to the decision maker simply by choosing a 0-differentially private mechanism that always outputs θ : $m(\theta|\theta') = 1$ for all $\theta' \in \Theta$. Thus, for the guarantee offered by differential privacy to be meaningful, the publication mechanism must be chosen *before* the designer learns the contents of the database — as we require in our model.

However, even if the designer does not observe the database before choosing a publication mechanism, she may still be better informed than an attacker about the data generating process from which the database is drawn. Her choice of mechanism may thus allow such an attacker to infer that information, and update his beliefs about an individual respondent’s type in a way that the mechanism — and hence the differential privacy criterion (1) — does not account for. We note that this is not a privacy violation in the sense considered by differential privacy, because the extra information learned *does not come from the database being protected*.²²

Like most of the formal privacy literature, we assume the true value of the population statistic enters directly into the payoffs of data users (decision makers). This means we abstract away from other sources of uncertainty — like measurement error and sampling variability — about the extent to which the undistorted data reflect some true underlying state of the world that data users actually care about. This is without loss of generality, if we interpret the data users’ payoffs as their expected utility conditional on the true value of the population statistic. Furthermore, the model assumes data users make decisions by observing the published statistic and then updating their beliefs based on that observation. In reality, data users sometimes treat data as though it is published without error.

²¹Note that, when it is known, dividing by the number of respondents N transforms the count in the category of interest into the *proportion* of the population in that category.

²²As Dwork et al. (2014) put it, “Differential privacy promises to protect individuals from any additional harm that they might face due to their data being in the private database x that they would not have faced had their data not been part of x . Although individuals may indeed face harm once the results $\mathcal{M}(x)$ of a differentially private mechanism \mathcal{M} have been released, differential privacy promises that the probability of harm was not significantly increased by their choice to participate.” (Dwork et al. (2014), p. 20)

3 Data Publication as Information Design

The differential privacy condition (1) can be reformulated as a restriction on the set of posterior beliefs that the mechanism can induce. Observe that if $\pi \in \Delta(\Theta)$ is the posterior belief of an agent who views realization s from the data publication mechanism (S, m) , Bayes' rule tells us that for any two databases θ and θ' that differ in a single entry,

$$\frac{\pi(\theta')}{\pi(\theta)} = \frac{m(s|\theta')\pi_0(\theta')}{\sum_{t \in \{0,1\}^N} m(s|t)\pi_0(t)} \bigg/ \frac{m(s|\theta)\pi_0(\theta)}{\sum_{t \in \{0,1\}^N} m(s|t)\pi_0(t)} = \frac{m(s|\theta')}{m(s|\theta)} \frac{\pi_0(\theta')}{\pi_0(\theta)}. \quad (3)$$

Hence, (S, m) is ϵ -differentially private if and only if each posterior belief π that it induces satisfies

$$\left| \log \left(\frac{\pi(\theta)}{\pi(\theta')} \right) - \log \left(\frac{\pi_0(\theta)}{\pi_0(\theta')} \right) \right| \leq \epsilon \text{ for each } \theta, \theta' \in \Theta \text{ with } \theta_{-n} = \theta'_{-n} \text{ for some } n. \quad (4)$$

We call the set of posteriors that satisfy (4) the ϵ -differentially private posteriors $K(\epsilon, \pi_0)$. In words, the ϵ -differentially private posteriors are those such that for any two databases θ and θ' that differ in a single entry, the difference in the log ratio of their probabilities under the prior and posterior is at most ϵ . As we show formally in the appendix (Lemma 6), this set is a closed convex polyhedron which does not intersect the edges of the probability simplex, and which contains the prior in its (relative) interior.²³ Figure 1 illustrates.




Figure 1: ϵ -differentially private posteriors. Consider a setting with binary data ($T = 1$), $N = 2$, $\epsilon = 1$, and $(\pi_0(0,0), \pi_0(0,1), \pi_0(1,0), \pi_0(1,1)) = (\frac{1}{3}, \frac{1}{6}, \frac{1}{6}, \frac{1}{3})$. Left panel: The blue shaded region of the probability simplex is described by the constraint $-\epsilon \leq \log(\pi((0,1))/\pi((0,0))) - \log(\pi_0((0,1))/\pi_0((0,0))) \leq \epsilon$ bounding the amount of information that an ϵ -differentially private mechanism can reveal about θ_2 when $\theta_1 = 0$. Middle panel: The region enclosed by the green lines is described by the differential privacy constraint on the amount of information revealed about θ_1 when $\theta_2 = 0$; its intersection with the region from the left panel is shaded in blue. Right panel: The blue shaded region is the intersection of the regions described by the constraints in (4), i.e., the set $K(\epsilon, \pi_0)$ of ϵ -differentially private posteriors. Note that $K(\epsilon, \pi_0)$ is asymmetric because there is no constraint on the ratio $\pi((0,0))/\pi((1,1))$, since the type profiles $(0,0)$ and $(1,1)$ are not adjacent.

By representing ϵ -differential privacy as a restriction to mechanisms that induce ϵ -differentially

²³It is worth noting that the robustness of ϵ -differential privacy to *post-processing*, i.e., the composition of the mechanism with a (possibly stochastic) map, arises from the convexity of the set of ϵ -differentially private posteriors: By Bayes' rule, any posterior induced by a signal that has been garbled by composition must be a convex combination of posteriors induced by the original signal.

private posteriors, we can succinctly write the designer's problem as a constrained information design problem. That is, the designer's problem simplifies to one of choosing a Bayes-plausible distribution $\nu \in \Delta(\Delta(\Theta))$ ²⁴ of posteriors (i.e., a distribution whose expected value is the prior) whose support is constrained to lie in the set $K(\epsilon, \pi_0)$.

Proposition 1 (Differentially Private Data Publication as Information Design). *The mechanism (S, m) solves the designer's problem (2) if and only if it induces a distribution of posteriors which solves*

$$\max_{\nu \in \Delta(K(\epsilon, \pi_0))} \{E_\nu v(\pi) \text{ s.t. } E_\nu \pi = \pi_0\}. \quad (5)$$

Since the designer's incentives are aligned with the decision maker's, his value function in (5) is convex. If the designer were not constrained by differential privacy, Kamenica and Gentzkow (2011) show that the solution to her problem would be straightforward: she should simply induce the posteriors at the vertices of $\Delta(\Theta)$ by publishing the true value of θ . With an ϵ -differentially private mechanism, this is impossible: she is restricted to inducing posteriors in $K(\epsilon, \pi_0)$. But convexity — or equivalently, Blackwell's (1953) theorem — still allows her to restrict attention to certain posteriors; namely, the extreme points (i.e., vertices) of the polyhedron $K(\epsilon, \pi_0)$.

This sharpens the conclusions that can be drawn from results in the information design literature when they are applied to the designer's problem. For any $K \subseteq \Delta(\Theta)$, denote the restriction of v to K as $v_K : K \rightarrow \mathbb{R} \cup \{-\infty\}$. Following Kamenica and Gentzkow (2011), define the K -restricted concavification $V_K : K \rightarrow \mathbb{R}$ of v_K as the smallest concave function that lies above the designer's value function v on the set of posteriors $K \subseteq \Delta(\Theta)$.²⁵

Corollary 1 (Characterization of Optimal Data Publication Mechanisms).

- i. *The maximized value of the designer's problem (2) is $V_{K(\epsilon, \pi_0)}(\pi_0)$.*
- ii. *There is a mechanism which solves the designer's problem (2) and induces a distribution of posteriors $\nu^* \in \Delta(\Delta(\Theta))$ such that*
 - (a) *Each $\pi \in \text{supp } \nu^*$ is an extreme point of $K(\epsilon, \pi_0)$;*
 - (b) *Each $\pi \in \text{supp } \nu^*$ achieves the privacy bound (4) between at least $(T+1)^N - 1$ pairs of databases: $\left| \log \left(\frac{\pi(\theta)}{\pi(\theta')} \right) - \log \left(\frac{\pi_0(\theta)}{\pi_0(\theta')} \right) \right| = \epsilon$ for at least $(T+1)^N - 1$ distinct combinations (θ, θ') that have $\theta_{-i} = \theta'_{-i}$ for some i ;*
 - (c) *The support of ν^* is a linearly independent set of vectors in $\mathbb{R}^{(T+1)^N}$; and thus²⁶*
 - (d) *ν^* is the unique Bayes-plausible distribution of posteriors with support $\text{supp } \nu^*$.*

Corollary 1 applies standard results from information design to give a characterization of optimal differentially private data publication that we use to establish our main results. Part (i) shows

²⁴Recall that since $\Delta(\Theta)$ is the set of posterior beliefs over databases, $\Delta(\Delta(\Theta))$ is the set of distributions of posteriors over databases.

²⁵Formally, let $V_K(\pi) \equiv \sup\{z | (\pi, z) \in \text{conv}(\text{Gr}(v_K))\}$, where $\text{Gr}(v_K) \equiv \{(\pi, v_K(\pi)) | \pi \in K\}$ denotes the graph of v_K .

²⁶For concreteness, let each belief $\pi \in \Delta(\Theta)$ be represented by the vector in $\mathbb{R}^{(T+1)^N}$ whose n th entry corresponds to the probability $\pi(\theta)$ it places on the database θ that is the binary number for n .

that the value of the problem has the familiar concavification characterization from Kamenica and Gentzkow (2011) — but because of differential privacy, this concavification takes place on the set of ϵ -differentially private posteriors, rather than the entire simplex. This places additional structure on the problem’s solution: Part (ii) places an upper bound on the number of posteriors induced by the mechanism — and hence the number of outputs it produces (iic) — and a lower bound on the number of privacy bounds that each of those posteriors must attain (iib). It also allows us to focus on the *set* of posteriors that the mechanism induces — which, without loss, is linearly independent (iic) — and ignore the probabilities with which it induces them (iid).

4 Conditionally Independent Noise and Oblivious Mechanisms

Section 3 shows that the designer’s problem amounts to choosing a Bayes-plausible distribution of posterior beliefs about the database. This problem is challenging in part because of its dimensionality: The space of posteriors about the database $\Delta(\Theta)$ has dimension $(T + 1)^N - 1$, the number of possible databases minus one.

But because the decision maker’s payoff only depends on the population statistic ω , his interim payoff only depends on his belief $\pi \in \Delta(\Theta)$ about the database through its *projection* onto the lower-dimensional space $\Delta(\Omega)$ of beliefs about the population statistic. That is, letting $P : \Delta(\Theta) \rightarrow \Delta(\Omega)$ be the projection operator defined by $P\pi(\omega) = \sum_{\theta: \omega_\theta = \omega} \pi(\theta)$, we have

$$v(\pi) \equiv \max_{a \in A} E_\pi u(a, \omega_\theta) = \max_{a \in A} E_{P\pi} u(a, \omega) \equiv \hat{v}(P\pi). \quad (6)$$

Figure 2 illustrates.




Figure 2: Projection onto $\Delta(\Omega)$ in Figure 1. A posterior belief $\pi \in \Delta(\Theta)$ about the database θ is projected onto the space $\Delta(\Omega)$ — shown here embedded in $\Delta(\Theta)$ — of beliefs about the population statistic ω .

Since the designer works to maximize the expected value of this payoff, we might then expect that she can restrict her attention to data publication mechanisms that depend only on the true value of the population statistic; i.e., those that add noise that is, conditional on ω , independent of the structure of the database θ . Following the formal privacy literature, we refer to such mechanisms as *oblivious*. This class includes most mechanisms used in practice in the kinds of settings we

consider. If oblivious mechanisms are without loss, she can solve her problem (5) by choosing a distribution on the set of posteriors about the population statistic $\Delta(\Omega)$ directly, rather than choosing a distribution of posteriors about the database and projecting them onto $\Delta(\Omega)$.

It turns out that the harmlessness of focusing on oblivious mechanisms depends crucially on the type of data that the designer has. In Theorem 1, we show that in settings with magnitude data, oblivious mechanisms are never without loss: There are always ϵ -differentially private mechanisms that induce distributions of posteriors about the population statistic ω that cannot be replicated using an oblivious mechanism. But with binary data, the optimality of oblivious mechanisms hinges on the structure of the prior π_0 . In particular, when respondents are anonymous, in the sense that the prior over databases is symmetric, we show in Theorem 2 that a designer can safely restrict attention to oblivious mechanisms. This assumption holds naturally in many settings, such as those where the data are i.i.d.

This has practical implications for the design of differentially private publication mechanisms: When firms and statistical agencies face the design problem considered in this paper, they mostly choose to adopt mechanisms that add conditionally independent noise to the true population statistic. As we show, selecting a mechanism from this class is optimal whenever the data is binary and the designer and data users view respondents as interchangeable, but not necessarily otherwise. With binary data, if the identity of a respondent carries information about his or her type, the designer may be able to exploit that information to design a mechanism that produces more useful output while maintaining differential privacy. And with magnitude data, the designer can exploit the fact that (unlike with binary data) differential privacy restricts beliefs differently at different databases with the same population statistic.

4.1 Characterizing Differential Privacy for Oblivious Mechanisms

Formally, we say a data publication mechanism (S, m) is *oblivious* if $m(\cdot|\theta) = m(\cdot|\theta')$ whenever $\omega_\theta = \omega_{\theta'}$.²⁷ Hence, there is a function $\sigma : \Omega \rightarrow \Delta(S)$ such that $m(\cdot|\theta) = \sigma(\cdot|\omega_\theta)$ for each $\theta \in \Theta$; we abuse notation and write (S, σ) to denote such a mechanism. The class of oblivious mechanisms includes most differentially private mechanisms used in practice to publish information about the count or proportion of individuals with a certain characteristic. In particular, the widely used *Gaussian* and *geometric* mechanisms are each oblivious, since they publish the sum of the true population statistic ω_θ and a random variable.

Because oblivious mechanisms only provide information about the population statistic, they are completely characterized by the distribution $\tau \in \Delta(\Delta(\Omega))$ of posterior beliefs that they induce about the population statistic, rather than the distribution of beliefs they induce about the database more generally. Such a distribution can be induced by an oblivious mechanism precisely when its expectation $E_\tau \mu$ is equal to the prior about the population statistic $\mu_0 \equiv P\pi_0$.

When a mechanism is oblivious, Proposition 2 shows that the differential privacy criterion (1) simplifies to a limit on the amount that *moving to an adjacent population statistic* (i.e., from $\omega - t$

²⁷In the terminology of the computer science literature on formal privacy, the population statistic ω_θ is the outcome of a *sum query* (with magnitude data) or *counting query* (with binary data) applied to θ , and our definition specifies that a mechanism is oblivious with respect to that query.

to ω , or vice versa, for $1 \leq t \leq T$) can change the distribution of the mechanism's realizations. Intuitively, if two values ω, ω' of the population statistic differ by no more than T , there are a pair of databases with those statistics which differ in only one entry and by exactly the difference $\omega - \omega'$. Consequently, differential privacy bounds the ratio of the posterior probabilities of those states (7).

Proposition 2 (Differential Privacy for Oblivious Mechanisms). *Suppose (S, σ) is an oblivious data publication mechanism. Then the following are equivalent:*

- i. (S, σ) is ϵ -differentially private.
- ii. $\left| \log \left(\frac{\sigma(s|\omega)}{\sigma(s|\omega-t)} \right) \right| \leq \epsilon$ for each $s \in S$, $\omega \in \{1, \dots, NT\}$, and $1 \leq t \leq \min\{T, \omega\}$.
- iii. For each posterior belief about the population statistic $\mu \in \Delta(\Omega)$ induced by (S, σ) ,
 $\left| \log \left(\frac{\mu(\omega)}{\mu(\omega-t)} \right) - \log \left(\frac{\mu_0(\omega)}{\mu_0(\omega-t)} \right) \right| \leq \epsilon$ for each $\omega \in \{1, \dots, NT\}$ and $1 \leq t \leq \min\{T, \omega\}$. (7)

We call the set of posterior beliefs about the population statistic that satisfy (7) the *oblivious ϵ -differentially private posteriors* $K_\Omega(\epsilon, \mu_0) \subset \Delta(\Omega)$. Figure 3 illustrates in the binary data setting described in Figure 1.




Figure 3: Oblivious ϵ -differentially private posteriors. Recall that in Figure 1, we had binary data ($T = 1$), $N = 2$, $\epsilon = 1$, and prior $(\pi_0(0,0), \pi_0(0,1), \pi_0(1,0), \pi_0(1,1)) = (\frac{1}{3}, \frac{1}{6}, \frac{1}{6}, \frac{1}{3})$. Hence, $\mu_0 = P\pi_0 = [\frac{1}{3} \ \frac{1}{3} \ \frac{1}{3}]'$. The region of the probability simplex enclosed by the blue dotted lines is described by the constraint $-\epsilon \leq \log(\mu(1)/\mu(0)) - \log(\mu_0(1)/\mu_0(0)) \leq \epsilon$; the region enclosed by the red dotted lines is described by the constraint $-\epsilon \leq \log(\mu(2)/\mu(1)) - \log(\mu_0(2)/\mu_0(1)) \leq \epsilon$; their intersection is $K_\Omega(\epsilon, \mu_0)$. Note that $K_\Omega(\epsilon, \mu_0)$ is asymmetric because there is no constraint on the ratio $\mu(0)/\mu(2)$, since states 0 and 2 are not adjacent.

4.2 When Are Oblivious Mechanisms Without Loss?

Proposition 2 reveals the key difference between differential privacy's restrictions on the information disclosed by an oblivious mechanism and its restrictions on the information disclosed *about the population statistic* by a non-oblivious mechanism. In general, the designer can induce a distribution τ of posterior beliefs about the population statistic whenever it is the projection onto $\Delta(\Omega)$ of a Bayes-plausible distribution on the set of ϵ -differentially private posteriors; i.e., whenever there exists a Bayes-plausible $\nu \in \Delta(K(\epsilon, \pi_0))$ such that $\tau(\mu) = \nu(P^{-1}(\mu))$ for each μ .²⁸ The set of distributions that satisfy this criterion coincides with the set of distributions that can be induced with an

²⁸That is, whenever it is the pushforward measure $\nu \circ P^{-1}$ of some Bayes-plausible $\nu \in \Delta(K(\epsilon, \pi_0))$ under the projection map P .

ϵ -differentially private *oblivious* mechanism precisely when

$$K_{\Omega}(\epsilon, \mu_0) = PK(\epsilon, \pi_0),$$

i.e., when the projection of any ϵ -differentially private posterior onto $\Delta(\Omega)$ is an oblivious ϵ -differentially private posterior (Lemma 9 in the appendix).




Figure 4: Projection of $K(\epsilon, \pi_0)$ onto $\Delta(\Omega)$ in Figures 1 and 3. Left panel: The set of ϵ -differentially private posteriors (blue) and the space $\Delta(\Omega)$ of posteriors about the population statistic (brown). Right panel: Since π_0 is symmetric and data is binary (i.e., $T = 1$) in Figures 1 and 3, the projection operator P carries $K(\epsilon, \pi_0)$ to the set of oblivious ϵ -differentially private posteriors (blue).

The main results of this section show that whether this is the case depends crucially on the kind of data (and hence, the kind of population statistic) that the designer wishes to publish information about. With magnitude data ($T > 1$), the answer is always no, regardless of the prior π or the privacy loss budget ϵ . Consequently, oblivious mechanisms are *never* without loss.

Theorem 1 (Oblivious Mechanisms and Magnitude Data). *With magnitude data, $K_{\Omega}(\epsilon, \mu_0) \neq PK(\epsilon, \pi_0)$: For any prior μ_0 and privacy loss budget ϵ , there exist ϵ -differentially private data publication mechanisms that induce distributions of beliefs about the population statistic that cannot be replicated with an oblivious mechanism that is also ϵ -differentially private.*

The intuition is as follows. Because an oblivious mechanism only depends on the database through the population statistic, viewing its output must change the log ratio of the probabilities of a pair of databases by the same amount whenever they have the same population statistics:

$$\left| \log \left(\frac{\pi(\theta)}{\pi(\theta')} \right) - \log \left(\frac{\pi_0(\theta)}{\pi_0(\theta')} \right) \right| = \left| \log \left(\frac{\pi(\hat{\theta})}{\pi(\hat{\theta}')} \right) - \log \left(\frac{\pi_0(\hat{\theta})}{\pi_0(\hat{\theta}')} \right) \right| \quad (8)$$

for each $\theta, \theta', \hat{\theta}, \hat{\theta}' \in \Theta$ with $\omega_{\theta} = \omega_{\hat{\theta}}$ and $\omega_{\theta'} = \omega_{\hat{\theta}'}$.

Since this change is the quantity bounded by differential privacy, requiring a mechanism to be oblivious effectively adds privacy bounds for pairs of databases θ, θ' that do not differ in one entry, but merely happen to have the same population statistics as some other pair $\hat{\theta}, \hat{\theta}'$ that does. The key to establishing Theorem 1 is showing that with magnitude data, these additional constraints never

become redundant when they are projected onto the space $\Delta(\Omega)$ of beliefs about the population statistic.

But with binary data ($T = 1$), it turns out that projection onto $\Delta(\Omega)$ *does* render these additional constraints redundant, so long as the prior distribution π_0 is symmetric: $\pi_0(\theta) = \pi_0(\theta')$ whenever θ is a permutation of θ' .²⁹ When the prior is symmetric in this way, we say that *respondents are anonymous*. Theorem 2 shows that when data is binary, anonymity ensures that $K_\Omega(\epsilon, \mu_0) = PK(\epsilon, \pi_0)$ — and hence that oblivious mechanisms are without loss.

Theorem 2 (Oblivious Mechanisms and Binary Data). *With binary data and anonymous respondents, $K_\Omega(\epsilon, \mu_0) = PK(\epsilon, \pi_0)$: For each ϵ -differentially private data publication mechanism, there exists an ϵ -differentially private oblivious data publication mechanism that induces the same distribution of posterior beliefs about the population statistic.*

As a consequence of Theorem 2, whenever respondents are anonymous and data is binary, the designer’s information design problem (5) can be simplified to one of choosing a distribution on the lower-dimensional space $\Delta(\Omega)$ of posterior beliefs about ω .

Corollary 2 (Differentially Private Data Publication as Information Design: Binary Data). *If respondents are anonymous and data is binary, the oblivious mechanism (S, σ) solves the designer’s problem (2) if and only if it induces a distribution of posteriors about the population statistic that solves*

$$\max_{\tau \in \Delta(K_\Omega(\epsilon, \mu_0))} \{E_\tau \hat{v}(\mu) \text{ s.t. } E_\tau \mu = \mu_0\}. \quad (9)$$

The role played by the respondents’ anonymity in Theorem 2 is subtle, as is the role of the kind of data that the designer collects about them. Because the respondents’ types take binary values with binary data, databases with the same population statistic must be permutations of each other.³⁰ Hence, each differs by a single entry from the same number of databases whose population statistic is lower by 1. Moreover, since these databases also differ from one another only by permutation, a symmetric prior places equal probability on each of them. As a consequence of these facts, the bounds that differential privacy places on the posterior likelihood ratios of databases that differ *in one entry* sum to the bounds it places on the posterior likelihood ratios of states that differ *by one* induced by an oblivious mechanism. That is, the constraints that characterize $K_\Omega(\epsilon, \mu_0)$ are the projections of the constraints that characterize $K(\epsilon, \pi_0)$, and so $K_\Omega(\epsilon, \mu_0) = PK(\epsilon, \pi_0)$.

We emphasize that anonymity leads to Theorem 2 *indirectly*. In particular, the prior’s symmetry is needed not because it eliminates differences between the *probabilities* of permutations of θ , but because it renders the differences between the sets of databases *adjacent* to those permutations irrelevant for differential privacy. Consequently, it is unnecessary when there are only two respondents: With $N = 2$, permuting θ does not change the set of databases that differ from θ in one entry.

Proposition 3 (Oblivious Mechanisms with Two Respondents). *With binary data and two respondents, $K_\Omega(\epsilon, \mu_0) = PK(\epsilon, \pi_0)$: For each ϵ -differentially private data publication mechanism (S, m) , there exists*

²⁹In other words, when respondent types $\{\theta_n\}_{n=1}^N$ are exchangeable random variables.

³⁰That is, if $\omega_\theta = \omega_{\theta'}$, then $\{t | \omega_t = \omega_\theta - 1 \text{ and } t_{-i} = \theta_{-i} \text{ for some } i\}$ and $\{t | \omega_t = \omega_{\theta'} - 1 \text{ and } t_{-i} = \theta'_{-i} \text{ for some } i\}$ have the same number of elements.

an ϵ -differentially private oblivious data publication mechanism (S, σ) that induces the same distribution of posterior beliefs about the population statistic.

4.3 Discussion

In general, the key reason for the stark contrast between the conclusions of Theorems 1 and 2 is that with binary data, databases with the same population statistic are identical up to permutation, whereas with magnitude data, they are not. This suggests that in settings with magnitude data, even though we cannot restrict attention to mechanisms that depend only on the population statistic, we may be able to restrict attention to mechanisms that are invariant under permutation. Formally, we say that a mechanism is *permutation-invariant* if $m(s|\theta) = m(s|\theta')$ whenever θ is a permutation of θ' . In the online appendix, we characterize the differential privacy criterion for permutation-invariant mechanisms, and prove a result (Proposition 4) showing that they are without loss whenever respondents are anonymous.

Proposition 4 (Permutation-Invariant Mechanisms). *If respondents are anonymous, then for each ϵ -differentially private data publication mechanism, there exists a permutation-invariant data publication mechanism that is also ϵ -differentially private and induces the same distribution of posterior beliefs about the population statistic.*

The conclusions of this section are markedly different from those of Ghosh et al. (2012), who consider a model with binary data. When the decision maker seeks to minimize a function of the distance between their action and the population statistic,³¹ they show that oblivious mechanisms are without loss if the decision maker has a prior belief μ_0 about the population statistic ω , but is ambiguity-averse, in the maxmin sense of Gilboa et al. (1989), and finds all beliefs π_0 about the database θ that are consistent with μ_0 (i.e., with $P\pi_0 = \mu_0$) to be plausible. In contrast, when decision makers are expected utility maximizers, we show that regardless of their preferences, the desirability of restricting attention to oblivious mechanisms depends both on the kind of data that the designer is publishing information about, and on whether certain respondents are *a priori* more likely than others to have the characteristic in question. We emphasize that such symmetry is crucial for our binary data result (Theorem 2): In Section S.1 of the Online Appendix, we give an example showing that when respondents are not anonymous, there may be a non-oblivious mechanism that outperforms every oblivious ϵ -differentially private mechanism.

Brenner and Nissim (2014) also provide negative results showing that optimality results for environments with binary data do not extend to the case of magnitude data. In particular, they show that, in contrast to the Ghosh et al. (2012) result for binary data, there is no mechanism, oblivious or otherwise, that is always optimal whenever data users wish to minimize some function of the distance between the true population statistic and (postprocessed) output. Theorem 1 makes a different contribution: There no *single* oblivious mechanism that is always optimal for the designer, regardless of the prior or the decision maker's preferences. In fact, is not always optimal for the designer to restrict attention to *any* oblivious mechanisms .

³¹That is, when $u(a, \omega) = \tilde{u}(|a - \omega|, \omega)$ for \tilde{u} nondecreasing in $|a - \omega|$.

Finally, we note that with binary data and anonymous respondents, Theorem 2 allows us to reduce Corollary 1 to a set of statements about objects in the lower-dimensional space $\Delta(\Omega)$ of posteriors about the population statistic, and sharpen its characterization of the designer's problem. For any $K \subseteq \Delta(\Omega)$, denote the restriction of \hat{v} to K as $\hat{v}_K : K \rightarrow \mathbb{R} \cup \{-\infty\}$, and the concavification of \hat{v}_K as $\hat{V}_K : K \rightarrow \mathbb{R}$. Then we have the following corollary to Theorem 2.

Corollary 3 (Characterization of Optimal Oblivious Mechanisms). *Suppose that respondents are anonymous and data is binary.*

- i. *The maximized value of the designer's problem (2) is $\hat{V}_{K_\Omega(\epsilon, \mu_0)}(\mu_0)$.*
- ii. *There is an oblivious mechanism which solves the designer's problem (2) and induces a distribution of posteriors about the population statistic $\tau^* \in \Delta(\Delta(\Omega))$ such that*
 - (a) *Each $\mu \in \text{supp } \tau^*$ is an extreme point of $K_\Omega(\epsilon, \mu_0)$;*
 - (b) *Each $\mu \in \text{supp } \tau^*$ achieves the privacy bound (7) at each value of the population statistic:*

$$\left| \log \left(\frac{\mu(\omega)}{\mu(\omega-1)} \right) - \log \left(\frac{\mu_0(\omega)}{\mu_0(\omega-1)} \right) \right| = \epsilon \text{ for all } \mu \in \text{supp } \tau^* \text{ and } \omega \in \Omega \setminus \{0\};$$
 - (c) *The support of τ^* is a linearly independent set of vectors in \mathbb{R}^{N+1} ; and thus³²*
 - (d) *τ^* is the unique Bayes-plausible distribution of posteriors with support $\text{supp } \tau^*$.*

5 Supermodular Payoffs and the Geometric Mechanism

In a large class of applications where data is binary and respondents are anonymous, we can move beyond Corollary 3's characterization. In particular, we focus on binary data settings where the decision maker's set of actions is totally ordered — and hence can be represented as a subset $A \subseteq \mathbb{R}$ of the real numbers — and he views higher actions and higher population statistics as complementary, in the sense that his payoff function is supermodular:

$$u(a', \omega') - u(a, \omega') \geq u(a', \omega) - u(a, \omega) \text{ for each } a' > a \text{ and } \omega' > \omega.$$

In this section, we show that these features — binary data, anonymous respondents, ordered actions, and supermodular payoffs — ensure that the well-known ϵ -geometric mechanism (Ghosh et al., 2012) described in the introduction always solves the designer's problem (Theorem 3).³³ Because of

³²Recall that with binary data ($T = 1$), we have $|\Omega| = |\{0, 1, \dots, N\}| = N + 1$.

³³As described in the introduction, the ϵ -geometric mechanism adds two-sided geometrically distributed noise to the true population statistic and publishes the result. Hence, it can be formally represented as $(\mathbb{Z}, \sigma_\epsilon^g)$, where $\sigma_\epsilon^g(s|\omega) = \left(\frac{1-e^{-\epsilon}}{1+e^{-\epsilon}} \right) e^{-\epsilon|s-\omega|}$. While the ϵ -geometric mechanism produces more outputs than there are states — which, by Corollary 3 (iic), is unnecessary — it induces the same distribution of posterior beliefs about the population statistic as the truncated ϵ -geometric mechanism $(\Omega, \hat{\sigma}_\epsilon^g)$ (Ghosh et al., 2012) with signal distribution

$$\hat{\sigma}_\epsilon^g(s|\omega) = \begin{cases} \left(\frac{1-e^{-\epsilon}}{1+e^{-\epsilon}} \right) e^{-\epsilon|s-\omega|}, & 0 < s < N, \\ \left(\frac{1}{1+e^{-\epsilon}} \right) e^{-\epsilon|s-\omega|}, & s \in \{0, N\}. \end{cases}$$

Hence, if a data provider wants to use as few outputs as possible, or just avoid publishing negative outputs, it can truncate the geometrically-distributed noise it adds without changing the mechanism's value to decision makers.

the nature of the information structures induced by ϵ -differentially private mechanisms, we cannot do so by appealing to dominance results from the literature (e.g., Quah and Strulovici (2009); Athey and Levin (2018)). Instead, we introduce a new notion of *Uniform-Peaked Relative Risk (UPRR) dominance* among information structures, and show that the ϵ -geometric mechanism is UPRR-dominant within the class of ϵ -differentially private data publication mechanisms. We then complete the argument by showing that a UPRR-dominant information structure is superior in any supermodular decision problem with actions on the real line (Theorem 4).

5.1 Optimality of the Geometric Mechanism

To explain why the geometric mechanism is optimal in these settings, we start by returning to the binary data setting of Figures 1-4, with $N = 2$, symmetric prior $\pi_0 = (\frac{1}{3}, \frac{1}{6}, \frac{1}{6}, \frac{1}{3})$, and hence $\mu_0 = P\pi_0 = [\frac{1}{3} \ \frac{1}{3} \ \frac{1}{3}]'$. There, our results from Section 4 tell us that when designing a publication mechanism,

- we need only consider oblivious mechanisms (Theorem 2);
- we can focus on the distribution of posteriors *about the population statistic* that the mechanism induces (Corollary 2); and
- it is without loss to consider such distributions supported by a linearly independent set of extreme points of $K_\Omega(\epsilon, \mu_0)$, and when we do, the support we choose pins down the distribution (Corollary 3).

Consequently, the designer only ever needs to consider two mechanisms in this setting: one oblivious mechanism that induces the set of posteriors in the left panel of Figure 5, and another that induces the set of posteriors in the right panel of Figure 5.




Figure 5: Possible solutions to the designer's problem (9) in Figures 1-4. In each panel, the extreme points of $K_\Omega(\epsilon, \mu_0)$ are labeled μ_w, μ_x, μ_y , and μ_z . Left panel: One possible solution is the unique Bayes-plausible distribution of posteriors about the population statistic with support $\{\mu_x, \mu_y, \mu_z\}$. Right panel: The alternative is the unique Bayes-plausible distribution with support $\{\mu_w, \mu_y\}$.

These distributions focus on providing information about two different aspects of the population statistic ω : the one in the left panel provides information about whether ω is high or low, while the one in the right panel focuses on providing information about whether ω is central or extreme.

Which distribution is better for the decision maker depends on which kind of information is more relevant to his choice of action.³⁴

For the class of problems that we focus on in this section — those where actions are totally ordered, and payoffs are supermodular — the answer is straightforward. When the decision maker's actions are totally ordered, they can be ranked from lowest to highest. When, in addition, his payoff function is supermodular, rightward shifts in his belief about the population statistic make higher actions more attractive relative to lower ones. Consequently, it is more valuable for him to know whether the population statistic is high or low than whether it is central or extreme — and so the distribution in the left panel is optimal.

This intuition generalizes to binary data settings with more than two respondents. Observe that the beliefs in the left panel of Figure 5 are precisely those that satisfy the upper privacy bound $\log\left(\frac{\mu(\omega)}{\mu(\omega-1)}\right) - \log\left(\frac{\mu_0(\omega)}{\mu_0(\omega-1)}\right) = \epsilon$ for each value of ω below some (posterior-specific) cutoff x , and the lower privacy bound $\log\left(\frac{\mu(\omega)}{\mu(\omega-1)}\right) - \log\left(\frac{\mu_0(\omega)}{\mu_0(\omega-1)}\right) = -\epsilon$ at each value of ω above the cutoff. That is, they are precisely those posteriors about the population statistic that achieve the upper privacy bound on *lower sets* of the form $(\Omega \setminus \{0\}) \cap (-\infty, x]$. As we show in Lemma 1, this property characterizes the posterior beliefs about ω that are induced by the ϵ -geometric mechanism.

Lemma 1 (Posteriors Produced by the Geometric Mechanism). *With binary data, the ϵ -geometric mechanism induces precisely those oblivious ϵ -differentially private posteriors that satisfy the upper privacy bound for $0 < \omega \leq x$ and the lower privacy bound for $\omega > x$ for some $x > 0$. That is, letting τ_ϵ^g denote the distribution of posteriors about the population statistic that is induced by the ϵ -geometric mechanism,*

$$\text{supp } \tau_\epsilon^g = \left\{ \mu \in \text{ext}(K_\Omega(\epsilon, \mu_0)) \mid \frac{\mu(\omega)/\mu(\omega-1)}{\mu_0(\omega)/\mu_0(\omega-1)} = \begin{cases} e^\epsilon, & \omega \leq x, \\ e^{-\epsilon}, & \omega > x, \end{cases} \text{ for some } x \in \mathbb{Z} \right\}.$$

To understand Lemma 1, observe that if a decision maker's belief about the population statistic after viewing the output x is μ , then for any $\omega > 0$, we have

$$\log\left(\frac{\mu(\omega)}{\mu(\omega-1)}\right) - \log\left(\frac{\mu_0(\omega)}{\mu_0(\omega-1)}\right) = \log\left(\frac{\sigma_\epsilon^g(x|\omega)}{\sigma_\epsilon^g(x|\omega-1)}\right) = \epsilon(|x - (\omega-1)| - |x - \omega|).$$

If $\omega \leq x$, then this simplifies to ϵ , and the upper privacy bound is satisfied; if $\omega > x$, then it simplifies to $-\epsilon$, and the lower privacy bound is satisfied.³⁵ Since they achieve the upper privacy bound for population statistics $\omega \leq x$ and the lower privacy bound for population statistics $\omega > x$, the posteriors induced by each output x of the geometric mechanism are collapsed around population statistic x (or the population statistic closest to x , if $x \notin [0, N]$) as much as the differential privacy constraint will allow.

This is intuitively optimal when the decision maker wants to minimize a nonincreasing function of the distance between their action and the true population statistic, as in the “postprocessing” ex-

³⁴Or, more precisely, more relevant to his marginal payoff from choosing one action over another.

³⁵Note that this means the posteriors induced by outputs $x \leq 0$ achieve all $N - 1$ lower privacy bounds, and so (since $\Delta(\Omega)$ has dimension $N - 1$) must be identical; likewise, each output $x \geq N$ induces the same posterior, which achieves all $N - 1$ upper privacy bounds.

ercise of Ghosh et al. (2012), since it allows the decision maker to be more certain that his chosen action is near the true value of ω . But equally intuitively, it is suboptimal when the decision maker is more interested in whether the population statistic is central or extreme — in which case a mechanism that induces the distribution in the right panel of Figure 5 (or an $N > 2$ analogue) is better.

Theorem 3 shows that the crucial feature that leads to the geometric mechanism’s optimality is not that the decision maker wants to *match* the population statistic (as in Ghosh et al. (2012)), but rather that he wants to take *higher actions when his belief is more skewed toward higher population statistics*. In particular, whenever the decision maker’s actions can be ordered from lowest to highest (e.g., in the introduction, from fewer to more buses), and the decision maker’s payoffs are supermodular — that is, his marginal benefit of taking a higher action is increasing in the population statistic — the ϵ -geometric mechanism outperforms any other ϵ -differentially private oblivious mechanism.

Theorem 3 (Optimality of the Geometric Mechanism for Supermodular Problems). *Suppose that respondents are anonymous and the data is binary. If the decision maker’s actions A are a compact set of real numbers, and the decision maker’s payoff function u is supermodular, then the ϵ -geometric mechanism solves the designer’s problem (2).*

5.2 UPRR-Dominant Information Structures

Theorem 3 cannot rely on existing results from the literature on comparing information structures. Oblivious mechanisms which induce extreme points of $K_\Omega(\epsilon, \mu_0)$ are not comparable in the Blackwell (1953) order: By definition, any extreme point of $K_\Omega(\epsilon, \mu_0)$ cannot be expressed as a convex combination of the others, and so no distribution supported by those points can be a garbling of another. Moreover, while the posterior beliefs induced by the ϵ -geometric mechanism are ordered by the monotone likelihood ratio property, the posteriors induced by many other oblivious mechanisms — such as one that induces the posteriors in the right panel of Figure 5 — need not be ranked in *any* meaningful sense (e.g., first-order stochastic dominance).³⁶ This makes results like those of Lehmann (1988), Quah and Strulovici (2009), Athey and Levin (2018), and Kim (2023) unavailable for our purposes, since they do not rank the information structures induced by the publication mechanisms we need to compare. Instead, we introduce a new notion of dominance among information structures in which the ϵ -geometric data publication mechanism dominates all other ϵ -differentially private oblivious data publication mechanisms. To emphasize the generality of this dominance criterion (and our results about it), we will often refer to ω as the *state* in this section.

Definition (Uniform-Peaked Relative Risk Dominance). For $\tau, \tau' \in \Delta(\Delta(\Omega))$, we say that τ dominates τ' in the sense of *uniform-peaked relative risk (UPRR)* and write $\tau \succeq_{UPRR} \tau'$ if each $\mu \in \text{supp } \tau$ has a *peak* $\omega^*(\mu) \in \Omega$ such that for each $\mu' \in \text{supp } \tau'$, $\mu(\omega)/\mu'(\omega)$ is nondecreasing on $(-\infty, \omega^*(\mu)]$ and nonincreasing on $[\omega^*(\mu), \infty)$.

³⁶Recall that a family of distributions $\{\mu_s\}_{s \in S}$ on a finite set $\Omega \subset \mathbb{R}$ is MLRP-ordered if $\mu_{s'}(\omega')/\mu_{s'}(\omega) \geq \mu_s(\omega')/\mu_s(\omega)$ for each $\omega' > \omega$ and $s' > s$. If $\{\mu_s\}_{s \in S}$ are the posterior beliefs induced by a signal $(S, \sigma : \Omega \rightarrow \Delta(S))$, they are MLRP-ordered whenever σ has the MLRP, i.e., whenever $\sigma(s'|\omega')/\sigma(s'|\omega) \geq \sigma(s|\omega')/\sigma(s|\omega)$ for each $\omega' > \omega$ and $s' > s$ (Quah and Strulovici, 2009; Milgrom, 1981).

In words, one distribution of posteriors on $\Omega \subset \mathbb{N}$ UPRR-dominates another if each posterior in the dominant distribution's support concentrates relatively more probability around one state — its peak — than any posterior in the dominated distribution's support. Equivalently, the *relative risk* $\mu(\omega)/\mu'(\omega)$ of a state ω under a posterior μ in the dominant distribution's support against some posterior μ' in the latter distribution is always increasing in the state below the first posterior's peak $\omega^*(\mu)$ and decreasing above it.³⁷ This dominance relation \succeq_{UPRR} is not a partial order or a preorder, but as Lemma 2 shows, it is transitive.³⁸

Lemma 2 (Transitivity of UPRR Dominance). *Suppose that $\tau, \tau', \tau'' \in \Delta(\Delta(\Omega))$ are each Bayes-plausible. If $\tau \succeq_{UPRR} \tau'$ and $\tau' \succeq_{UPRR} \tau''$, then $\tau \succeq_{UPRR} \tau''$.*

Example 2. UPRR dominance allows us to compare the two distributions of posteriors described in Figure 5. The distribution τ_b from the right panel of Figure 5 is supported by the two extreme points of $K_\Omega(\epsilon, \mu_0)$ that achieve exactly one upper privacy bound (at either $\omega = 1$ or $\omega = 2$). The distribution τ_ϵ^δ from the left panel of Figure 5, on the other hand, induces the extreme points of $K_\Omega(\epsilon, \mu_0)$ that achieve the upper bound for both states, neither state, and for state 1 only, respectively.

Below, in Figure 6, we graph all four posterior beliefs $\mu \in \text{ext}(K_\Omega(\epsilon, \mu_0))$ that achieve the privacy bound at each value of ω , as well as the relative risk $\mu(\omega)/\mu'(\omega)$ for each $\mu \in \text{supp } \tau_\epsilon^\delta$ and each $\mu' \in \text{supp } \tau_b$. One can see that for each $\mu \in \text{supp } \tau_\epsilon^\delta$, the relative risk against either posterior $\mu' \in \text{supp } \tau_b$ has the same peak $\omega^*(\mu)$: when μ attains both lower privacy bounds, $\omega^*(\mu) = 0$; when μ attains the upper privacy bound for $\omega = 1$ and the lower privacy bound for $\omega = 2$, $\omega^*(\mu) = 1$; and when μ attains both upper privacy bounds, $\omega^*(\mu) = 2$. Hence, $\tau_\epsilon^\delta \succeq_{UPRR} \tau_b$.

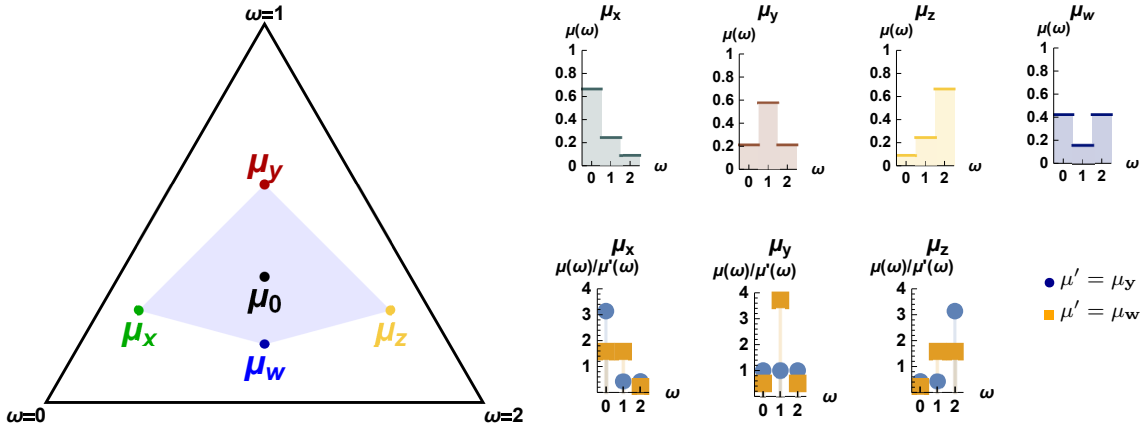


Figure 6: UPRR dominance in Figure 5. Left panel: extreme points of $K_\Omega(\epsilon, \mu_0)$ in the simplex $\Delta(\Omega)$. Top right panel: probability mass functions of extreme points of $K_\Omega(\epsilon, \mu_0)$. Bottom right panel: Relative risk $\frac{\mu(\omega)}{\mu'(\omega)}$ under beliefs $\mu \in \text{supp } \tau_\epsilon^\delta$ versus beliefs $\mu' \in \text{supp } \tau_b$.

³⁷Note that for Bayes-plausible distributions, the UPRR order can be equivalently defined in terms of the signals that generate those distributions: If τ and τ' are induced by the data publication mechanisms (S, σ) and (S', σ') , then $\tau \succeq_{UPRR} \tau'$ (and so we can write $\sigma \succeq_{UPRR} \sigma'$) if and only if each $s \in S$ has a peak $\hat{\omega}^*(s)$ such that for each $s' \in S'$, $\sigma(s|\omega)/\sigma'(s'|\omega)$ is nondecreasing on $(-\infty, \hat{\omega}^*(s)]$ and nonincreasing on $[\hat{\omega}^*(s), \infty)$.

³⁸The UPRR relation fails to be either reflexive (consider a distribution that induces one posterior that is single-peaked and another that is uniform) or irreflexive, as is necessary for a strict preorder (as Lemma 3 shows, the distribution of posteriors τ_ϵ^δ UPRR-dominates itself).

Lemma 3 shows that the UPRR dominance of the ϵ -geometric mechanism demonstrated in Example 2 extends to *all* ϵ -differentially private data publication mechanisms with any number of respondents.

Lemma 3 (Geometric Mechanisms are UPRR-Dominant). *If $\tau \in \Delta(\Delta(\Omega))$ is induced by an ϵ -differentially private oblivious mechanism, then $\tau_\epsilon^g \succeq_{\text{UPRR}} \tau$, where for each $\mu \in \text{supp } \tau_\epsilon^g$, $\omega^*(\mu)$ is given by the largest $\omega \in \Omega$ for which $\frac{\mu(\omega)/\mu(\omega-1)}{\mu_0(\omega)/\mu_0(\omega-1)} = e^\epsilon$, or 0 if no such ω exist.*

For intuition, recall that each posterior $\mu \in \Delta(\Omega)$ induced by the geometric mechanism satisfies the upper privacy bound at every value of ω at or below some x and the lower privacy bound at each higher value of ω . Hence, these posteriors each concentrate probability around their peak $\omega^*(\mu) = x$ as much as the differential privacy constraint (7) will allow. This is precisely what is necessary for the distribution they support to UPRR-dominate the distributions of posteriors about the population statistic induced by every other differentially private oblivious mechanism.

We conclude with our general comparative statics result on UPRR dominance. Theorem 4 shows that among information structures which induce finitely many posteriors, those that UPRR-dominant are more useful for supermodular decision problems with actions on the real line. Since the distribution of posterior beliefs about the state induced by the geometric mechanism UPRR-dominates every other distribution of oblivious ϵ -differentially private posteriors (Lemma 3), and the designer's problem can be solved by inducing a distribution of posteriors with finite support (Corollary 3 (iic)), Theorem 3 follows.

Theorem 4 (UPRR Dominance Implies Dominance in Supermodular Problems). *If $\tau, \tau' \in \Delta(\Delta(\Omega))$ are Bayes-plausible and have finite support, and $\tau \succeq_{\text{UPRR}} \tau'$, then for any compact $A \subseteq \mathbb{R}$ and continuous supermodular function $h : \Omega \times A \rightarrow \mathbb{R}$, $E_\tau[\max_{a \in A} E_\mu[h(\omega, a)]] \geq E_{\tau'}[\max_{a \in A} E_\mu[h(\omega, a)]]$.*

Observe that when a decision maker's payoffs are supermodular, the costs associated with making a "mistake" and choosing an action that is optimal in state ω^* instead of one that is optimal in the true state ω are increasing in the distance between ω and ω^* . Intuitively, then, if two beliefs μ and μ' induce the same action, but μ places relatively more mass than μ' on states closer to the state $\omega^*(\mu)$ where that action is optimal, the expected costs of these "mistakes" should be lower under μ than under μ' , no matter how quickly they are increasing in the distance $|\omega - \omega^*(\mu)|$. If this were true for every pair of beliefs μ and μ' induced, respectively, by the Bayes-plausible distributions τ and τ' — as is the case when τ UPRR-dominates τ' — it would suggest that any decision maker with supermodular payoffs would be better off under the information structure τ than under τ' .

Theorem 4 confirms this conjecture. Its argument — which we discuss in detail in Appendix A — relies on what we call a *Frechét representation* of the information structures in question. This tool allows us to represent a distribution of posterior beliefs as the joint distribution of the state and (a smoothed version of) the cumulative distribution function of some function of the belief. The key insight leading to Theorem 4 is that whenever one information structure UPRR-dominates another, then there is a particular Frechét representation of the former that dominates *any* Frechét representation of the latter in the supermodular stochastic order (Proposition 5).³⁹

³⁹For a thorough discussion of the supermodular stochastic order, see Shaked and Shanthikumar (2007).

6 Conclusion

Our analysis introduces the tools of information design to the problem of differentially private data publication. Here, we offer a summary of the practical implications of our results. First, Corollaries 1 and 3 show how to select an optimal publication mechanism using concavification results from the information design literature. Even if a data provider does not wish to derive an optimal mechanism explicitly, these results establish criteria for checking whether any given mechanism is potentially optimal by checking the number of privacy bounds that it attains. Second, most differentially private data publication mechanisms used in practice are *oblivious*, in the sense that they add conditionally independent noise to the true value of the population statistic that data users are interested in, rather than publishing output in a way that depends more generally on the database. We show that using a mechanism from this class is never without loss with magnitude data (Theorem 1), but always optimal with binary data in the very common case where respondents are anonymous (Theorem 2) — for instance, if the data are i.i.d. Finally, we show that among oblivious mechanisms, the commonly used *geometric mechanism* is optimal whenever data users view their actions and the population statistic of interest as complementary, in the sense that their payoffs are supermodular. We do so by using a novel result that compares the value of information structures in supermodular decision problems.

Outside of the binary, supermodular case considered by Theorem 3, the mechanisms that solve the designer’s problem are not necessarily easy for humans to interpret. For instance, the optimal mechanism in Example 1 is a vector with tens of thousands of parameters. We thus have very little to say qualitatively about what they might look like. But the practical message of our paper for settings with magnitude data, or where data users face non-supermodular problems, is not a set of intuitions about when one mechanism is better than another. Rather, it is that there are (potentially significant) benefits to constructing a publication mechanism by computing the solution to an optimization problem — which, like other information design problems, can be formulated as a linear program — rather than using an off-the-shelf mechanism. With binary data and anonymous respondents, we can greatly reduce the dimensionality of that problem so that its solution can be computed more quickly (Theorem 2), but with magnitude data, we cannot (Theorem 1; Example 1).

While we focus on the standard differential privacy criterion, other variants of differential privacy exist. A popular one is *approximate* or (ϵ, δ) -*differential privacy*, which adds a tolerance δ to the constraint on the proportion by which the probability of a signal may differ conditional on databases that differ in one entry.⁴⁰ In subsequent work, Pan et al. (2025) show that Theorem 2 extends to this concept: When data is binary and respondents are anonymous, oblivious mechanisms are without loss among (ϵ, δ) -differentially private mechanisms. However, they also show that unlike ϵ -differential privacy, (ϵ, δ) -differential privacy does not amount to a constraint on the set of posterior beliefs that the designer can induce, and so the designer faces a constrained information design problem more complex than (5).⁴¹

⁴⁰Formally, (S, m) is (ϵ, δ) -differentially private if $m(s|\theta) \leq e^\epsilon m(s|\theta') + \delta$ for any databases $\theta, \theta' \in \Theta$ with $\theta_{-i} = \theta'_{-i}$ for some i .

⁴¹Pan et al. (2025) also generalize our results by considering misalignment between the preferences of the data users

Several applications suggest generalizations of our model. First, in many settings, data users interact with one another, instead of making independent decisions. For instance, when considering whether to attend a gathering, individuals may not only consider the information published by the designer about COVID-19 prevalence in their community, but also the attendance decisions of others. This suggests extending the model to allow for such interactions among the decision makers. In particular, it may be worthwhile to identify the circumstances under which Theorem 3 continues to hold, i.e., those in which the geometric mechanism remains optimal in applications where data users have supermodular decision problems that are also affected by the actions taken by others.

Second, many data providers wish to publish information about multiple characteristics, each of which might affect the data user’s decision problem through the number of respondents that have it. Such *multidimensional* data publication problems have been considered by Brenner and Nissim (2014). They give a counterexample showing that in contrast to the one-dimensional setting, the geometric mechanism is not universally optimal for all data users who minimize symmetric loss functions; that is, the result of Ghosh et al. (2012) no longer holds in such settings. To our knowledge, no characterization of optimal publication mechanisms in multidimensional settings has appeared in the privacy literature. While it is straightforward to show that a characterization result analogous to Corollary 1 can be given in such settings, the circumstances under which oblivious mechanisms are without loss are less clear. We leave this question to future work.

References

- ABOWD, J. M., G. L. BENEDETTO, S. L. GARFINKEL, S. A. DAHL, A. N. DAJANI, M. GRAHAM, M. B. HAWES, V. KARWA, D. KIFER, H. KIM, P. LECLERC, A. MACHANAVAJJHALA, J. P. REITER, R. RODRIGUEZ, I. M. SCHMUTTE, W. N. SEXTON, P. E. SINGER, AND L. VILHUBER (2020): “The modernization of statistical disclosure limitation at the US Census Bureau,” in *Washington, DC: US Census Bureau*. Available at: <https://www2.census.gov/adrm/CED/Papers/CY20/2020-08-AbowdBenedettoGarfinkelDahletal-The%20modernization%20of.pdf>.
- ABOWD, J. M. AND I. M. SCHMUTTE (2019): “An Economic Analysis of Privacy Protection and Statistical Accuracy as Social Choices,” *American Economic Review*, 109, 171–202.
- ABOWD, J. M., I. M. SCHMUTTE, W. N. SEXTON, AND L. VILHUBER (2019): “Why the Economics Profession Must Actively Participate in the Privacy Protection Debate,” *AEA Papers and Proceedings*, 109, 397–402.
- ACQUISTI, A., C. TAYLOR, AND L. WAGMAN (2016): “The Economics of Privacy,” *Journal of economic Literature*, 54, 442–492.
- ALI, S. N., G. LEWIS, S. VASSERMAN, ET AL. (2022): “Voluntary Disclosure and Personalized Pricing,” .
- APPLE INC. (2017): “Apple Differential Privacy Technical Overview,” Retrieved 2022-2-8, https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf.
- ARRIETA-IBARRA, I., L. GOFF, D. JIMÉNEZ-HERNÁNDEZ, J. LANIER, AND E. G. WEYL (2018): “Should We Treat Data as Labor? Moving Beyond “Free”,” *AEA Papers and Proceedings*, 108, 38–42.
- ATHEY, S. AND J. LEVIN (2018): “The Value of Information in Monotone Decision Problems,” *Research in Economics*, 72, 101–116.

and those of the designer.

- BLACKWELL, D. (1953): “Equivalent Comparisons of Experiments,” *The Annals of Mathematical Statistics*, 265–272.
- BRENNER, H. AND K. NISSIM (2014): “Impossibility of Differentially Private Universally Optimal Mechanisms,” *SIAM Journal on Computing*, 43, 1513–1540.
- CHETTY, R. AND J. N. FRIEDMAN (2019): “A Practical Method to Reduce Privacy Loss when Disclosing Statistics Based on Small Samples,” Working Paper 25626, National Bureau of Economic Research.
- DINUR, I. AND K. NISSIM (2003): “Revealing information while preserving privacy,” in *Proceedings of the Twenty-second ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, New York, NY, USA: ACM, PODS ’03, 202–210.
- DOVAL, L. AND V. SKRETA (2021): “Constrained Information Design: Toolkit,” *arXiv preprint arXiv:1811.03588*.
- DOW, A., A. HERDAĞDELEN, P. MOHASSEL, A. POMPE, AND B. STATE (2020): “Protecting privacy in Facebook mobility data during the COVID-19 response,” Retrieved 2020-12-30. <https://research.fb.com/blog/2020/06/protecting-privacy-in-facebook-mobility-data-during-the-covid-19-response/>.
- DWORK, C., F. MCSHERRY, K. NISSIM, AND A. SMITH (2006): “Calibrating Noise to Sensitivity in Private Data Analysis,” in *Theory of Cryptography Conference*, Springer, 265–284.
- DWORK, C., A. ROTH, ET AL. (2014): “The Algorithmic Foundations of Differential Privacy.” *Found. Trends Theor. Comput. Sci.*, 9, 211–407.
- ECHENIQUE, F. AND K. HE (2021): “Screening p -Hackers: Dissemination Noise as Bait,” *arXiv preprint arXiv:2103.09164*.
- EILAT, R., K. ELIAZ, AND X. MU (2021): “Bayesian Privacy,” *Theoretical Economics*, 16, 1557–1603.
- EPSTEIN, L. G. AND S. M. TANNY (1980): “Increasing Generalized Correlation: A Definition and Some Economic Consequences,” *Canadian Journal of Economics*, 16–34.
- FIORETTO, F., P. V. HENTENRYCK, AND J. ZIANI (2024): “Differential Privacy Overview and Fundamental Techniques,” .
- FOOTE, A., J. K. HAHN, S. TIBBETS, AND L. WARREN (2021): “Post-Secondary Employment Outcomes (PSEO),” Tech. rep., U.S. Census Bureau, retrieved 5 Feb 2022, <https://lehd.ces.census.gov/doc/PSEOTechnicalDocumentation.pdf>.
- GENG, Q. AND P. VISWANATH (2015): “The Optimal Noise-Adding Mechanism in Differential Privacy,” *IEEE Transactions on Information Theory*, 62, 925–951.
- GHOSH, A., T. ROUGHGARDEN, AND M. SUNDARARAJAN (2012): “Universally Utility-maximizing Privacy Mechanisms,” *SIAM Journal on Computing*, 41, 1673–1693.
- GILBOA, I., D. SCHMEIDLER, ET AL. (1989): “Maxmin Expected Utility with Non-Unique Prior,” *Journal of Mathematical Economics*, 18, 141–153.
- GUEVARA, M. (2019): “Enabling developers and organizations to use differential privacy,” Retrieved 2020-12-30, <https://developers.googleblog.com/2019/09/enabling-developers-and-organizations.html>.
- HE, K., F. SANDOMIRSKIY, AND O. TAMUZ (2023): “Private private information,” *arXiv preprint arXiv:2112.14356*.
- HIDIR, S. AND N. VELLODI (2021): “Privacy, Personalization, and Price Discrimination,” *Journal of the European Economic Association*, 19, 1342–1363.
- HSU, J., M. GABOARDI, A. HAEBERLEN, S. KHANNA, A. NARAYAN, B. C. PIERCE, AND A. ROTH (2014): “Differential Privacy: An Economic Method for Choosing Epsilon,” *2014 IEEE 27th Computer Security Foundations Symposium*, 398–410.
- ICHIHASHI, S. (2020): “Online Privacy and Information Disclosure by Consumers,” *American Economic Review*, 110, 569–95.

- JONES, C. I. AND C. TONETTI (2020): “Nonrivalry and the Economics of Data,” *American Economic Review*, 110, 2819–58.
- KAMENICA, E. (2019): “Bayesian Persuasion and Information Design,” *Annual Review of Economics*, 11, 249–272.
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- KIFER, D. AND A. MACHANAVAJJHALA (2011): “No Free Lunch in Data Privacy,” in *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data*, New York, NY, USA: Association for Computing Machinery, SIGMOD ’11, 193–204.
- KIM, Y. (2023): “Comparing Information in General Monotone Decision Problems,” *Journal of Economic Theory*, 211, 105679.
- KOUFOGIANNIS, F., S. HAN, AND G. J. PAPPAS (2015): “Optimality of the Laplace Mechanism in Differential Privacy,” *arXiv preprint arXiv:1504.00065*, 10.
- LE TREUST, M. AND T. TOMALA (2019): “Persuasion with Limited Communication Capacity,” *Journal of Economic Theory*, 184, 104940.
- LEHMANN, E. (1988): “Comparing Location Experiments,” *The Annals of Statistics*, 521–533.
- MATYSKOVÁ, L. (2019): “Bayesian Persuasion With Costly Information Acquisition,” Working paper.
- MEYER, M. A. AND B. H. STRULOVICI (2013): “The Supermodular Stochastic Ordering,” CEPR Discussion Paper No. DP9486.
- MILGROM, P. R. (1981): “Good News and Bad News: Representation Theorems and Applications,” *The Bell Journal of Economics*, 380–391.
- NEAR, J. (2018): “Differential Privacy at Scale: Uber and Berkeley Collaboration,” in *Enigma 2018 (Enigma 2018)*, Santa Clara, CA: USENIX Association.
- PAI, M. M. AND A. ROTH (2013): “Privacy and Mechanism Design,” *ACM SIGecom Exchanges*, 12, 8–29.
- PAN, Y., Z. S. WU, H. XU, AND S. ZHENG (2025): “Differentially Private Bayesian Persuasion,” in *Proceedings of the ACM on Web Conference 2025*, 1425–1440.
- PERSICO, N. (1996): “Information Acquisition in Affiliated Decision Problems,” Working paper.
- QUAH, J. K.-H. AND B. STRULOVICI (2009): “Comparative Statics, Informativeness, and the Interval Dominance Order,” *Econometrica*, 77, 1949–1992.
- ROCKAFELLAR, R. T. (1997): *Convex Analysis*, Princeton Paperbacks, Princeton University Press.
- ROGERS, R., S. SUBRAMANIAM, S. PENG, D. DURFEE, S. LEE, S. K. KANCHI, S. SAHAY, AND P. AHAMMAD (2020): “LinkedIn’s Audience Engagements API: A Privacy Preserving Data Analytics System at Scale,” .
- SHAKED, M. AND J. G. SHANTHIKUMAR (2007): *Stochastic Orders*, Springer Science & Business Media.
- STRACK, P. AND K. H. YANG (2024): “Privacy Preserving Signals,” *Available at SSRN 4467608*.
- TANEVA, I. (2019): “Information Design,” *American Economic Journal: Microeconomics*, 11, 151–185.
- US CENSUS BUREAU (2021): “2020 Census Redistricting Data (Public Law 94-171) Summary File Development & Release Timeline,” Retrieved 2022-2-8, <https://www2.census.gov/programs-surveys/decennial/2020/program-management/data-product-planning/disclosure-avoidance-system/das-redistricting-development-timeline.pdf>.
- VILHUBER, L. AND I. M. SCHMUTTE (2016): “Proceedings from the 2016 NSF–Sloan Workshop on Practical Privacy,” .
- WOLFRAM RESEARCH INC. (2019): *Mathematica, Version 12.0*, Champaign, IL.
- YODER, N. (2022): “Designing Incentives for Heterogeneous Researchers,” *Journal of Political Economy*, 130, 2018–2054.

Appendix

A Construction of Theorem 4

In this appendix, we provide a more detailed description of the arguments underlying our general comparative statics result on the UPRR order, Theorem 4. These arguments proceed in three phases. First, we show that each Bayes-plausible distribution of posteriors with finite support can be represented by an element of the *Frechét class* $\mathcal{M}(\mu_0, U([0, 1]))$ — the class of cumulative distribution functions with marginal distributions μ_0 and $U([0, 1])$.^{42,43} Next, in Lemma 4, we describe the sense in which this representation is equivalent to the original distribution. Finally, Proposition 5 shows that a UPRR-dominant distribution has a representation which is dominant in the supermodular stochastic order.

To begin, suppose that $\tau \in \Delta(\Delta(\Omega))$ is a Bayes-plausible distribution with finite support, and label the posteriors in its support $\{\mu_j\}_{j=1}^J$. Then it is straightforward that we can represent τ as a random vector $(\omega, R_\tau(j))$, where R_τ is the cumulative distribution of the posterior's index. This random vector is only supported on $\Omega \times \{R_\tau(j)\}_{j=1}^J$. However, we can also represent τ as its “uniform smoothing” $(W, X) \sim F$, where, letting $Q_\tau(z) \equiv \inf\{x | R_\tau(x) \geq z\}$ denote the quantile function of j ,

$$F(w, x) = \int_0^x \sum_{y=0}^w f(y, z) dz, \text{ for “mixed density” } f(y, z) = \mu_j(y)|_{j=Q_\tau(z)},$$

since (W, X) is equal in distribution to $(\omega, R_\tau(j))$ on the latter vector's support.⁴⁴ Instead of concentrating probability on a finite set of x -values, (W, X) distributes probability between those points uniformly, conditional on W .

We call (W, X) the *Frechét representation* of τ with respect to the index function $\mathcal{J} : \text{supp } \tau \rightarrow \mathbb{R}$ defined by $\mathcal{J}(\mu_j) \equiv j$. We use this name precisely because F is a member of the Frechét class $\mathcal{M}(\mu_0, U([0, 1]))$: The marginal distribution of W is μ_0 because τ is Bayes-plausible, while the marginal distribution of X is $U([0, 1])$ because the probabilities that a posterior μ places on states in Ω must sum to one. Example 3 illustrates the Frechét representation concept in the context of Figures 1-6.

Example 3 (Frechét Representations of the Geometric Mechanism). Consider the ϵ -geometric mechanism $(\mathbb{Z}, \sigma_\epsilon^g)$ in the simple setting from Figures 1-6. This data publication mechanism induces the three extreme points of $K_\Omega(\epsilon, \mu_0)$ characterized by attaining the upper privacy bound at the sets $\Phi_\mu = \emptyset$, $\Phi_\mu = \{1\}$, and $\Phi_\mu = \{1, 2\}$, which are each displayed in the lower-left of Figure 7, and does so according to the distribution σ_ϵ^g displayed in the upper-left of Figure 7. By Lemma 3, these posteriors have peaks $\omega^*(\mu)$ of 0, 1, and 2, respectively. Then when the posteriors are indexed by $\omega^*(\mu)$, the Frechét representation of the distribution τ_ϵ^g has the mixed density shown on the right of Figure 7.

More generally, we can consider a distribution's Frechét representation not only with respect to an index function, but any real-valued function on the distribution's support. Formally, for any $\tau \in \Delta(\Delta(\Omega))$ with finite support, and any $t : \text{supp } \tau \rightarrow \mathbb{R}$, let $r_{\tau,t}(x) \equiv \tau(\{\mu | t(\mu) = x\}) = \tau(t^{-1}(x))$ be the probability mass function of $T = t(\mu)$; $R_{\tau,t}(x) \equiv \tau(\{\mu | t(\mu) \leq x\}) = \sum_{z \leq x} r_{\tau,t}(z)$ be its cumulative distribution function; and $Q_{\tau,t}(z) \equiv \inf\{x \in t(\text{supp } \tau) | R_{\tau,t}(x) \geq z\}$ be its quantile function. We say that the random vector (W, X) is the Frechét representation of τ with respect to t if its cumulative distribution function is $F(w, x) =$

⁴²In fact, one can show that $\Delta(\Delta(\Omega))$ is *equivalent* to $\mathcal{M}(\mu_0, U([0, 1]))$, in the sense that every joint distribution whose cdf lies in the latter set represents a distribution of posteriors in the former.

⁴³Recall that a distribution of posteriors τ is Bayes-plausible if $E_\tau \mu = \mu_0$.

⁴⁴We use the term “mixed density” for f since it is neither a probability mass function nor a probability density function: Conditional on W , X is a continuous random variable, while conditional on X , W is a discrete random variable.




Figure 7: Generating the Frechét representation of τ_ϵ^g with respect to the peak ω^* in Figures 1-6. Upper-left panel: The distribution τ_ϵ^g . Lower-left panel: The posterior beliefs in $\text{supp } \tau_\epsilon^g$. Right panel: The mixed density of the Frechét representation of τ_ϵ^g with respect to ω^* .

$\int_0^x \sum_{y=0}^w f(y, z) dz$, with mixed density $f(y, z)$ given by

$$f(y, z) = \sum_{\mu \in t^{-1}(Q_{\tau, t}(z))} \frac{\tau(\mu) \mu(y)}{r_{\tau, t}(Q_{\tau, t}(z))}.$$

In words, $f(y, z)$ is the weighted (by τ) average of the probability placed on state y by posteriors μ whose value under t , $t(\mu)$, is equal to the t -quantile of z , $Q_{\tau, t}(z)$. In particular, when t is one-to-one, $f(y, z)$ is the probability placed on state y by the unique posterior μ such that $t(\mu) = Q_{\tau, t}(z)$.

Example 3, continued. Consider once more the geometric mechanism from Figures 1-6. This time, let $t : \text{supp } \tau_\epsilon^g \rightarrow \mathbb{R}$ map the posterior that attains the upper privacy bound at both states 1 and 2 to the same value as the extreme point of $K_\Omega(\epsilon, \mu_0)$ which attains the upper bound at state 1 alone:

$$t(\mu) = \begin{cases} 0, & \frac{\mu(\omega)/\mu(\omega-1)}{\mu_0(\omega)/\mu_0(\omega-1)} = e^{-\epsilon} \text{ for each } \omega \in \{1, 2\}, \\ 1, & \text{otherwise.} \end{cases}$$

The Frechét representation of τ_ϵ^g with respect to t (Figure 8) “coarsens” the representation of τ_ϵ^g with respect to ω^* . Once again, it represents the uniform smoothing of the joint distribution of the state ω and the cdf of a function of the posterior belief μ , but now, the function maps two posteriors to the same value.




Figure 8: The Frechét representation of τ_ϵ^g with respect to t in Figures 1-6. For each $\omega \in \Omega$, representing τ_ϵ^g with respect to t instead of ω^* averages the representation’s mixed density $f(w, x)$ over the region of $\{\omega\} \times [0, 1]$ associated with posteriors that attain one or both upper privacy bounds.

The Frechét representation we define here is similar to the representation employed by Athey and Levin (2018). In order to prove their main result, they represent a real-valued signal using the joint distribution of its cdf (or some other order-preserving function which maps to the unit interval) and the state. Hence, their representation preserves the exogenously given ordering on signal realizations — and thus posterior beliefs — inherent in their setting. In contrast, the Frechét representation can accommodate multiple orderings of the posteriors induced by a data publication mechanism by varying the function t with respect to which the representation is taken. This flexibility is crucial to our results: In order to show that UPRR-dominance implies dominance in supermodular problems (Theorem 4), we employ a Frechét representation with respect to an endogenous function.⁴⁵

Lemma 4 shows that a Frechét representation is indeed a member of the Frechét class $\mathcal{M}(\mu_0, U([0, 1]))$, and makes precise the sense in which it is equivalent to the distribution of posteriors that it represents.

Lemma 4 (Equivalence of Frechét Representations). *Suppose that $\tau \in \Delta(\Delta(\Omega))$ is Bayes-plausible and has finite support, let $A \subseteq \mathbb{R}$ be compact, let (W, X) be the Frechét representation of τ with respect to $t : \text{supp } \tau \rightarrow A$, and let F be the cdf of (W, X) .*

- i. $F \in \mathcal{M}(\mu_0, U([0, 1]))$.
- ii. For any measurable function $h : \Omega \times A \rightarrow \mathbb{R}$, $E_\tau[E_\mu[h(\omega, t(\mu))]] = E[h(W, Q_{\tau,t}(X))]$.
- iii. For any upper semicontinuous function $h : \Omega \times A \rightarrow \mathbb{R}$, any $\tau' \in \Delta(\Delta(\Omega))$ with finite support, and any $s : \text{supp } \tau' \rightarrow A$,

$$E_\tau \left[\max_{a \in A} E_\mu [h(\omega, a)] \right] \geq E[h(W, Q_{\tau',s}(X))].$$

Parts (ii) and (iii) of Lemma 4 show two different facets of the Frechét representation. The first is a straightforward equivalence: Suppose we want to take the expectation under τ of a function h (such as a decision maker’s utility function) which depends on the state ω and some function t of the posterior μ (such as a selection from that decision maker’s optimal action correspondence). Part (ii) says we can do so by taking an expectation of a function of the Frechét representation (W, X) of τ with respect to t instead. In particular, if we let h take the arguments W (the “state” part of the Frechét representation) and $Q_{\tau,t}(X)$ (the t -quantile of the “cumulative distribution” part of the Frechét representation, X), then its expectation is identical to the expectation under τ that we care about.

The second is a dominance argument that does not require equivalence between the function t used in the Frechét representation and the function s whose quantile appears in the expectation on the smaller side of the inequality. There, instead of letting h take the second argument $Q_{\tau,t}(X)$ — the X th percentile value of $t(\mu)$ when μ is distributed according to τ — h ’s second argument is instead $Q_{\tau',s}(X)$, the X th percentile value of $s(\mu)$ when μ is distributed according to τ' . By necessity, this choice of its second argument must lead h to take a weakly lower expected value than the maximizer of its conditional expectation with respect to X would.⁴⁶ The proof then shows that choosing a to maximize $E[h(W, a)|X]$ at each $X \in [0, 1]$ cannot outperform choosing a to maximize $E_\mu[h(\omega, a)]$ at each $\mu \in \text{supp } \tau$. (In fact, it may do strictly worse, since — as in Figure 8 — the Frechét representation (W, X) may discard information from τ by pooling together posteriors which t maps to the same value.)

The Frechét representation is useful for our purposes because it allows us to rank distributions of posteriors using the *supermodular stochastic order* (e.g., Shaked and Shanthikumar (2007); Meyer and Strulovici (2013)). Formally, if \mathbf{X} and \mathbf{Y} are random vectors whose i th elements X_i and Y_i take values in a compact set

⁴⁵Specifically, we use a Frechét representation with respect to a selection from the decision maker’s solution correspondence.

⁴⁶That is, for $a^*(X) \in \arg \max_{a \in A} E[h(W, a)|X]$, we have $E[h(W, a^*(X))|X] \geq E[h(W, Q_{\tau',s}(X))|X]$, and so by the law of iterated expectations, $E[h(W, a^*(X))] \geq E[h(W, Q_{\tau',s}(X))]$.

$S_i \subseteq \mathbb{R}$, and for any supermodular function $h : \prod S_i \rightarrow \mathbb{R}$, we have $E[h(\mathbf{X})] \geq E[h(\mathbf{Y})]$, then we say \mathbf{X} dominates \mathbf{Y} in the supermodular stochastic order and write $\mathbf{X} \succeq_{SPM} \mathbf{Y}$. The supermodular stochastic order can only rank random vectors which belong to the same Frechét class, i.e., which have the same marginal distributions. When applied to the Frechét class $\mathcal{M}(\mu_0, U([0, 1]))$, Proposition 5 shows that the supermodular order's ranking of Frechét representations is concordant with the UPRR order's ranking of information structures, in the sense that the Frechét representation of a UPRR-dominant distribution of posteriors with respect to the peak ω^* SPM-dominates any Frechét representation of the UPRR-dominated distribution. Combining this result with Lemma 4 then yields Theorem 4.

Proposition 5 (UPRR-Dominance Implies Supermodular Stochastic Dominance). *Suppose that $\tau, \tau' \in \Delta(\Delta(\Omega))$ are Bayes-plausible and have finite support, and that $\tau \succeq_{UPRR} \tau'$ with peaks given by the function $\omega^* : \text{supp } \tau \rightarrow \mathbb{R}$. Then the Frechét representation (W, X) of τ with respect to ω^* dominates the Frechét representation (\hat{W}, \hat{X}) of τ' with respect to $t : \text{supp } \tau' \rightarrow \mathbb{R}$ in the supermodular stochastic order: for any supermodular function $h : \Omega \times [0, 1] \rightarrow \mathbb{R}$, $E[h(W, X)] \geq E[h(\hat{W}, \hat{X})]$.*

As is well known (see, e.g., Shaked and Shanthikumar (2007)), one bivariate random vector dominates another in the supermodular stochastic order if and only if the cumulative distribution function of the former vector lies above that of the latter. Hence, denoting the cdfs of the Frechét representations (W, X) and (\hat{W}, \hat{X}) in Proposition 5 by F and G , respectively, $(W, X) \succeq_{SPM} (\hat{W}, \hat{X})$ is equivalent to the statement that $F(w, x) \geq G(w, x)$ for each $w \in \Omega$ and $x \in [0, 1]$. Furthermore, because F and G are both elements of $\mathcal{M}(\mu_0, U([0, 1]))$, the latter statement is equivalent to a similar one concerning the Frechét representations' survival functions $\bar{F}(w, x) \equiv P(W > w \text{ and } X > x)$ and $\bar{G}(w, x) \equiv P(\hat{W} > w \text{ and } \hat{X} > x)$. In particular, $F(w, x) \geq G(w, x)$ for each $w \in \Omega$ and $x \in [0, 1]$ if and only if $\bar{F}(w, x) \geq \bar{G}(w, x)$ for each $w \in \Omega$ and $x \in [0, 1]$.

Intuitively, this means that (W, X) dominates (\hat{W}, \hat{X}) in the supermodular stochastic order if it is more tightly distributed around a nondecreasing path between $(0, 0)$ and $(N, 1)$. Proposition 5 shows that this must be true when $\tau \succeq_{UPRR} \tau'$: At every point x in the unit interval, the distribution of W conditional on $X = x$ is a convex combination of posteriors $\mu \in \text{supp } \tau$ whose peaks $\omega^*(\mu)$ are each given by the ω^* -quantile of x , $Q_{\tau, \omega^*}(x)$. Moreover, since each of these posteriors' relative risk $\mu(\omega)/\mu'(\omega)$ against a posterior μ' induced by the UPRR-dominated distribution τ' has a single peak at $Q_{\tau, \omega^*}(x)$, the same must be true of the convex combination's relative risk against those posteriors. Likewise, the distribution of \hat{W} conditional on $\hat{X} = x$ is a convex combination of posteriors $\mu' \in \text{supp } \tau'$; it follows that the likelihood ratio of the conditional distributions has a single peak at $Q_{\tau, \omega^*}(x)$. Consequently, at any fixed x , (W, X) must be more concentrated around $Q_{\tau, \omega^*}(x)$ than (\hat{W}, \hat{X}) . Proposition 5 shows that this concentration extends across all of $[0, 1]$ — i.e., that (W, X) is more concentrated around the graph of Q_{τ, ω^*} than (\hat{W}, \hat{X}) — and hence (W, X) dominates (\hat{W}, \hat{X}) in the supermodular stochastic order.

B Proofs⁴⁷

Characterization of the Designer's Problem and (Obviously) Differentially Private Posteriors

Lemma 5. *The distribution $\nu \in \Delta(\Delta(\Theta))$ of posterior beliefs can be induced with an ϵ -differentially private mechanism (S, m) if and only if ν is Bayes-plausible and $\text{supp } \nu \subset K(\epsilon, \pi_0)$.*

Proof. By Kamenica and Gentzkow (2011), ν can be induced with some mechanism (S, m) if and only if it is Bayes-plausible. From Bayes' rule (3) and the definition of differential privacy (1), that mechanism is ϵ -differentially private if and only if $\text{supp } \nu \subset K(\epsilon, \pi_0)$ as well. \square

Proof of Proposition 1 (Differentially Private Data Publication as Information Design) Follows immediately from Lemma 5. \square

Lemma 6 (Characterization of Differentially Private Posteriors). *For each $\epsilon > 0$,*

- i. $K(\epsilon, \pi_0)$ lies in the relative interior of $\Delta(\Theta)$.
- ii. π_0 lies in the relative interior of $K(\epsilon, \pi_0)$.
- iii. $K(\epsilon, \pi_0)$ is a closed convex polyhedron in $\mathbb{R}^{2^{N(T+1)}}$.

Proof. (i): If $\pi \in \text{rbd}(\Delta(\Theta))$, then $\pi(\theta) = 0$ for some $\theta \in \Theta$. We cannot have $\pi(\theta) = 0$ for all $\theta \in \Theta$, so it follows that there must exist $\theta', \theta'' \in \Theta$, with $\theta''_{-i} = \theta'_{-i}$ for some i , such that $\pi(\theta') = 0$ and $\pi(\theta'') > 0$. Then $|\log(\pi(\theta'')/\pi(\theta'))| = \infty$, so $\pi \notin K(\epsilon, \pi_0)$.

(ii): Since $\log(x)$ is a continuous function on $(0, \infty)$, for each $\theta, \theta' \in \Theta$, $\log(\pi(\theta)/\pi(\theta'))$ is continuous in π on $\text{ri}(\Delta(\Theta))$. Then for each $\theta, \theta' \in \Theta$ with $\theta_{-i} = \theta'_{-i}$ for some i , the set

$$\left\{ \pi \in \text{ri}(\Delta(\Theta)) \mid \log\left(\frac{\pi(\theta)}{\pi(\theta')}\right) \in \left(\log\left(\frac{\pi_0(\theta)}{\pi_0(\theta')}\right) - \epsilon, \log\left(\frac{\pi_0(\theta)}{\pi_0(\theta')}\right) + \epsilon\right) \right\}$$

is open in $\text{ri}(\Delta(\Theta))$ and hence $\text{aff}(\Delta(\Theta))$. Then so is their intersection,

$$\hat{K} = \left\{ \pi \in \text{ri}(\Delta(\Theta)) \mid -\epsilon < \log\left(\frac{\pi(\theta)}{\pi(\theta')}\right) - \log\left(\frac{\pi_0(\theta)}{\pi_0(\theta')}\right) < \epsilon \forall \theta, \theta' \text{ s.t. } \exists i, \theta_{-i} = \theta'_{-i} \right\}.$$

Then since \hat{K} is open and $\pi_0 \in \hat{K} \subset K(\pi_0, \epsilon)$, we have $\pi_0 \in \text{ri}(K(\pi_0, \epsilon))$, as desired.

(iii): We have

$$\begin{aligned} K(\epsilon, \pi_0) &= \left\{ \pi \in \Delta(\Theta) \mid e^{-\epsilon} \leq \frac{\pi(\theta)/\pi_0(\theta)}{\pi(\theta')/\pi_0(\theta')} \leq e^\epsilon \forall \theta, \theta' \text{ s.t. } \exists i, \theta_{-i} = \theta'_{-i} \right\} \\ &= \left\{ \pi \in \Delta(\Theta) \mid e^{-\epsilon} \pi(\theta') \frac{\pi_0(\theta)}{\pi_0(\theta')} - \pi(\theta) \leq 0 \leq e^\epsilon \pi(\theta') \frac{\pi_0(\theta)}{\pi_0(\theta')} - \pi(\theta) \forall \theta, \theta' \text{ s.t. } \exists i, \theta_{-i} = \theta'_{-i} \right\} \end{aligned}$$

⁴⁷For a set S , denote its relative interior by $\text{ri}(S)$ and its relative boundary by $\text{rbd}(S)$. For $S \subset \mathbb{Z}_+$, we write $\mathbf{1}_S$ to denote the indicator vector for S , i.e., the vector (of conformable dimension) with 1 in its i th entry if $i \in S$ and 0 in its i th entry otherwise. Likewise, we write $\mathbf{1}_n$ to denote the indicator vector for $n \in \mathbb{Z}_+$, i.e., $\mathbf{1}_{\{n\}}$. Further, write $\mathbf{1}$ to denote a (conformable) vector of ones. For vectors in \mathbb{R}^{L+1} , we adopt the convention of *zero-indexing* throughout the paper: that is, we number the entries of vectors starting from 0 rather than starting from 1, so that the entries range from 0 to L rather than 1 to $L+1$. Hence, we can represent each belief $\mu \in \Delta(\Omega)$ as the vector in \mathbb{R}^{NT+1} whose n th entry corresponds to the probability $\mu(n)$ it places on $\omega = n$; i.e., the indices of the vector's entries and the population statistics they represent are the same. Likewise, we can let each belief $\pi \in \Delta(\Theta)$ be represented by the vector in $\mathbb{R}^{2^{N(T+1)}}$ whose n th entry corresponds to the probability $\pi(\theta)$ it places on the database θ that is the T -ary number for n . (For simplicity, we write $\mathbf{1}_\theta \in \mathbb{R}^{2^{N(T+1)}}$ for the indicator vector for the index whose T -ary number is θ .)

Thus, $K(\epsilon, \pi_0)$ is the intersection of finitely many half-spaces described by the inequalities

$$\begin{aligned} \pi \cdot \left(\mathbf{1}_\theta - e^{-\epsilon \frac{\pi_0(\theta)}{\pi_0(\theta')}} \mathbf{1}_{\theta'} \right) &\geq 0, & \pi \cdot \left(e^{\epsilon \frac{\pi_0(\theta)}{\pi_0(\theta')}} \mathbf{1}_{\theta'} - \mathbf{1}_\theta \right) &\geq 0, & \forall \theta, \theta' \text{ s.t. } \exists i, \theta_{-i} = \theta'_{-i}; \\ \pi \cdot \mathbf{1} &\geq 1; & \pi \cdot \mathbf{1} &\leq 1. \end{aligned} \quad (10)$$

(From (i), this intersection lies inside the intersection of the half-spaces $\pi(\theta) \geq 0$, and so the latter can be omitted.) Then by definition, $K(\epsilon, \pi_0)$ is a closed convex polyhedron. \square

Corollary 4 (Characterization of Oblivious Differentially Private Posteriors). *For each $\epsilon > 0$,*

- i. $K_\Omega(\epsilon, \mu_0)$ lies in the relative interior of $\Delta(\Omega)$.
- ii. μ_0 lies in the relative interior of $K_\Omega(\epsilon, \mu_0)$.
- iii. $K_\Omega(\epsilon, \mu_0)$ is a closed convex polyhedron in \mathbb{R}^{N+1} .

Proof. Follows identically to Lemma 6. \square

Proof of Proposition 2 (Differential Privacy for Oblivious Mechanisms) ((i) \Rightarrow (ii)) For each $\omega \in \Omega \setminus \{0\}$, choose $i \in \{1, \dots, N\}$ and $t \in \{1, \dots, \min\{\omega, T\}\}$, and choose θ such that $\theta_i = t$ and $\omega_\theta = \omega$, and θ' such that $\theta_{-i} = \theta'_{-i}$ and $\theta_i = 0$. Then since (S, σ) is ϵ -differentially private, for each $s \in S$, $|\log(\sigma(s|\omega)/\sigma(s|\omega - t))| = |\log(\sigma(s|\omega_\theta)/\sigma(s|\omega_{\theta'}))| \leq \epsilon$, since $\omega_{\theta'} = \omega - t$; (ii) follows.

((ii) \Rightarrow (i)) If $\theta, \theta' \in \{0, \dots, T\}^N$ are such that $\theta_{-i} = \theta'_{-i}$ for some i , then either $\theta = \theta'$, in which case (1) holds trivially, or $|\omega_\theta - \omega_{\theta'}| = t$ for some $1 \leq t \leq \min\{T, \max\{\omega_\theta, \omega_{\theta'}\}\}$, in which case (ii) implies that for each $s \in S$, $|\log(\sigma(s|\omega_\theta)/\sigma(s|\omega_{\theta'}))| = |\log(\sigma(s|\omega_{\theta'})/\sigma(s|\omega_\theta))| \leq \epsilon$, and hence, since (S, σ) is oblivious, (1).

((ii) \Leftrightarrow (iii)) Follows from Bayes' rule, since

$$\frac{\mu(\omega')}{\mu(\omega)} = \frac{\sigma(s|\omega')\mu_0(\omega')}{\sum_{x \in \Omega} \sigma(s|x)\mu_0(x)} \bigg/ \frac{\sigma(s|\omega)\mu_0(\omega)}{\sum_{x \in \Omega} \sigma(s|x)\mu_0(x)} = \frac{\sigma(s|\omega')}{\sigma(s|\omega)} \frac{\mu_0(\omega')}{\mu_0(\omega)}.$$

Lemma 7 (Characterization of Extreme Points of $K(\epsilon, \pi_0)$).

- i. If the support of a distribution ν of ϵ -differentially private posteriors contains a belief that is not an extreme point of $K(\epsilon, \pi_0)$, then there is some distribution of posterior beliefs $\hat{\nu}$ with $\text{supp } \hat{\nu} \subseteq \text{ext}(K(\epsilon, \pi_0))$ that is more informative than ν (in the Blackwell sense).
- ii. Each extreme point of $K(\epsilon, \pi_0)$ attains at least $2^{N(T+1)} - 1$ privacy bounds: For each $\pi \in \text{ext}(K(\epsilon, \pi_0))$, $\left| \log \left(\frac{\pi(\theta)}{\pi(\theta')} \right) - \log \left(\frac{\pi_0(\theta)}{\pi_0(\theta')} \right) \right| = \epsilon$ for at least $2^{N(T+1)} - 1$ distinct combinations (θ, θ') that have $\theta_{-i} = \theta'_{-i}$ for some i .

Proof. (i): By Lemma 6 (iii), $K(\epsilon, \pi_0)$ is a convex closed polyhedron; since it is a subset of the simplex $\Delta(\Theta)$, it is bounded and thus compact. Then by the Minkowski-Weyl theorem, $K(\epsilon, \pi_0)$ has finitely many extreme points and $K(\epsilon, \pi_0) = \text{conv}(\text{ext}(K(\epsilon, \pi_0)))$. Then for each $\pi \in \text{supp } \nu$, we can write $\pi = \sum_{\pi' \in \text{ext}(K(\epsilon, \pi_0))} \lambda(\pi, \pi') \pi'$ for some $\{\lambda(\pi, \pi')\}_{\pi' \in \text{ext}(K(\epsilon, \pi_0))} \subset [0, 1]$ with $\sum_{\pi' \in \text{ext}(K(\epsilon, \pi_0))} \lambda(\pi, \pi') = 1$. Then λ is a mean-preserving stochastic transformation (in the sense of Blackwell (1953)); consequently, by Blackwell (1953) Theorem 2, $\hat{\nu}$ defined by $\hat{\nu}(\pi') = \int_{\Delta(\Omega)} \lambda(\pi, \pi') d\tau(\pi)$ is more informative than ν .

(ii): Recall that the extreme points of a polyhedron $P \subset \mathbb{R}^L$ defined by $P = \{x | x \cdot z_j \geq b_j, j \in \{1, \dots, d\}\}$ for $\{z_i\}_{i=1}^d \subset \mathbb{R}^L$ and $\{b_j\}_{j=1}^d \subset \mathbb{R}$ are precisely the basic feasible solutions of a linear program on P ; i.e., those $x \in P$ such that $x \cdot z_j = b_j$ for a linearly independent subset of the z_j with L elements. By Lemma 6 (iii), $K(\epsilon, \pi_0) \subset \mathbb{R}^{2^{N(T+1)}}$ is a closed convex polyhedron described by the inequalities in (10). It follows that at

least $2^{N(T+1)}$ of those inequalities must bind at any $\pi \in K(\epsilon, \pi_0)$. For any $\theta \in \Theta$ and $\pi \in K(\epsilon, \pi_0)$, $\pi(\theta) \neq 0$ by Lemma 6 (i), and so for any $\theta, \theta' \in \Theta$, we cannot have both

$$\pi \cdot \left(\mathbf{1}_\theta - e^{-\epsilon} \frac{\pi_0(\theta)}{\pi_0(\theta')} \mathbf{1}_{\theta'} \right) = 0 \quad \text{and} \quad \pi \cdot \left(e^\epsilon \frac{\pi_0(\theta)}{\pi_0(\theta')} \mathbf{1}_{\omega-1} - \mathbf{1}_\omega \right) = 0. \quad (11)$$

And since both $\pi \cdot \mathbf{1} \geq 1$ and $\pi \cdot (-\mathbf{1}) \geq -1$ must bind, and $\mathbf{1}$ and $-\mathbf{1}$ are linearly dependent, it follows that we must have $\left| \log \left(\frac{\pi(\theta)}{\pi(\theta')} \right) - \log \left(\frac{\pi_0(\theta)}{\pi_0(\theta')} \right) \right| = \epsilon$ for at least $2^{N(T+1)} - 1$ distinct combinations (θ, θ') that have $\theta_{-i} = \theta'_{-i}$ for some i . \square

Corollary 5 (Characterization of Extreme Points of $K_\Omega(\epsilon, \mu_0)$). *Suppose that data is binary.*

- i. *If the support of a distribution τ of oblivious ϵ -differentially private posteriors contains a belief that is not an extreme point of $K_\Omega(\epsilon, \mu_0)$, then there is some distribution of posterior beliefs about the population statistic $\hat{\tau} \in \Delta(\Delta(\Omega))$ with $\text{supp } \hat{\tau} \subseteq \text{ext}(K_\Omega(\epsilon, \mu_0))$ that is more informative than τ (in the Blackwell sense).*
- ii. *The extreme points of $K_\Omega(\epsilon, \mu_0)$ are precisely those for which the privacy bound is attained at all ω : $\text{ext}(K_\Omega(\epsilon, \mu_0)) = \left\{ \mu \in \Delta(\Omega) \mid \left| \log \left(\frac{\mu(\omega)}{\mu(\omega-1)} \right) - \log \left(\frac{\mu_0(\omega)}{\mu_0(\omega-1)} \right) \right| = \epsilon \forall \omega \in \Omega \setminus \{0\} \right\}$.*

Proof. (i): Follows identically to Lemma 7 (i).

(ii): Recall that the extreme points of a polyhedron $P \subset \mathbb{R}^{N+1}$ defined by $P = \{\mu \mid \mu \cdot z_i \geq b_i, i \in \{1, \dots, d\}\}$ for $\{z_i\}_{i=1}^d \subset \mathbb{R}^{N+1}$ and $\{b_i\}_{i=1}^d \subset \mathbb{R}$ are precisely the basic feasible solutions of a linear program on P ; i.e., those $\mu \in P$ such that $\mu \cdot z_i = b_i$ for a linearly independent subset of the z_i with $N+1$ elements. By Lemma 6 (iii), $K_\Omega(\epsilon, \mu_0)$ is a closed convex polyhedron described by the inequalities

$$\begin{aligned} \mu \cdot \left(\mathbf{1}_\omega - e^{-\epsilon} \frac{\mu_0(\omega)}{\mu_0(\omega-1)} \mathbf{1}_{\omega-1} \right) &\geq 0, & \mu \cdot \left(e^\epsilon \frac{\mu_0(\omega)}{\mu_0(\omega-1)} \mathbf{1}_{\omega-1} - \mathbf{1}_\omega \right) &\geq 0, & \omega \in \Omega \setminus \{0\}; \\ \mu \cdot \mathbf{1} &\geq 1; & \mu \cdot \mathbf{1} &\leq 1. \end{aligned} \quad (12)$$

Now by definition, $\mu \cdot \mathbf{1} = 1$ for each $\mu \in K_\Omega(\epsilon, \mu_0)$. For any $\omega \in \Omega$ and $\mu \in K_\Omega(\epsilon, \mu_0)$, $\mu(\omega) \neq 0$ by Lemma 6 (i), and so for any $\omega \in \Omega \setminus \{0\}$, we cannot have both

$$\mu \cdot \left(\mathbf{1}_\omega - e^{-\epsilon} \frac{\mu_0(\omega)}{\mu_0(\omega-1)} \mathbf{1}_{\omega-1} \right) = 0 \quad \text{and} \quad \mu \cdot \left(e^\epsilon \frac{\mu_0(\omega)}{\mu_0(\omega-1)} \mathbf{1}_{\omega-1} - \mathbf{1}_\omega \right) = 0. \quad (13)$$

Hence, if $\mu \in \text{ext}(K_\Omega(\epsilon, \mu_0))$, then for each $\omega \in \Omega \setminus \{0\}$, exactly one equation in (13) must hold to reach a total of $N+1$ linearly independent binding inequalities from (12); it follows that $|\log(\mu(\omega)/\mu(\omega-1)) - \log(\mu_0(\omega)/\mu_0(\omega-1))| = \epsilon$ for each $\omega \in \Omega \setminus \{0\}$.

Conversely, if $|\log(\mu(\omega)/\mu(\omega-1)) - \log(\mu_0(\omega)/\mu_0(\omega-1))| = \epsilon$ for each $\omega \in \Omega \setminus \{0\}$, then for each $\omega \in \Omega \setminus \{0\}$, one of the equations in (13) must hold; since the set of vectors $Z = \{\mathbf{1}\} \cup \{e^\epsilon \frac{\mu_0(\omega)}{\mu_0(\omega-1)} \mathbf{1}_{\omega-1} - \mathbf{1}_\omega \mid \omega \in S\} \cup \{e^{-\epsilon} \frac{\mu_0(\omega)}{\mu_0(\omega-1)} \mathbf{1}_{\omega-1} - \mathbf{1}_\omega \mid \omega \in (\Omega \setminus \{0\}) \setminus S\}$ is linearly independent for each $S \subseteq \Omega \setminus \{0\}$, it follows that μ is a basic feasible solution for a linear program on $K_\Omega(\epsilon, \mu_0)$ and hence $\mu \in \text{ext}(K_\Omega(\epsilon, \mu_0))$. \square

Lemma 8 (Distributions of Posteriors with Affine and Linear Independent Supports).

- i. *If $v \in \Delta(\Delta(\Theta))$ is Bayes-plausible and $\text{supp } v$ is finite,⁴⁸ there is a Bayes-plausible $v' \in \Delta(\Delta(\Theta))$ such that the posteriors in $\text{supp } v'$ are affinely independent and $\text{supp } v' \subseteq \text{supp } v$.*
- ii. *If $\{\pi_j\}_{j=1}^J \subset \Delta(\Theta)$ is affinely independent, it is linearly independent.*
- iii. *If $v \in \Delta(\Delta(\Theta))$ is Bayes-plausible and $\text{supp } v$ is linearly independent, then $\text{supp } v' \neq \text{supp } v$ for any Bayes-plausible $v' \in \Delta(\Delta(\Theta))$, $v' \neq v$.*

⁴⁸The finiteness assumption is unnecessary here, but is adopted for simplicity.

Proof. (i): Since $\text{supp } \nu$ is finite and ν is Bayes-plausible, we have $\sum_{\pi \in \text{supp } \nu} \nu(\pi) \pi = \pi_0$, and so $\pi_0 \in \text{conv}(\text{supp } \nu)$. Then by Carathéodory's theorem (e.g., Rockafellar (1997) Theorem 17.1) there is some affinely independent $\{\pi_j\}_{j=1}^J \subseteq \text{supp } \nu$ such that $\pi_0 \in \text{conv}(\{\pi_j\}_{j=1}^J)$. Then there exist $\{\lambda_j\}_{j=1}^J \subset [0, 1]$ such that $\sum_{j=1}^J \lambda_j = 1$ and $\sum_{j=1}^J \lambda_j \pi_j = \pi_0$. The claim follows by letting $\nu'(\pi_j) = \lambda_j$ for each $j \in \{1, \dots, J\}$.

(ii): We prove the contrapositive. Suppose that $\{\pi_j\}_{j=1}^J$ is linearly dependent. Then there exist $\{\lambda_j\}_{j=1}^J \subset \mathbb{R}$ such that $\sum_{j=1}^J \lambda_j \pi_j = 0$ and $\lambda_j \neq 0$ for some j . Then $0 = \sum_{\theta \in \Theta} \sum_{j=1}^J \lambda_j \pi_j(\theta) = \sum_{j=1}^J \lambda_j \sum_{\theta \in \Theta} \pi_j(\theta)$. Since $\{\mu_j\}_{j=1}^J \subset \Delta(\Theta)$, $\sum_{\theta \in \Theta} \pi_j(\theta) = 1$ for each $j \in \{1, \dots, J\}$. It follows that $0 = \sum_{j=1}^J \lambda_j$ and thus $\{\mu_j\}_{j=1}^J$ is affinely dependent.

(iii): We proceed by contradiction: Suppose that $\text{supp } \nu' = \text{supp } \nu$ for some Bayes-plausible $\nu' \neq \nu$. Since $\text{supp } \nu$ is linearly independent, it must be finite. Since ν and ν' are both Bayes-plausible, we must have $\sum_{\pi \in \text{supp } \nu} \nu(\pi) \pi = \sum_{\pi \in \text{supp } \nu} \nu'(\pi) \pi = \pi_0$. Then we have $\sum_{\pi \in \text{supp } \nu} (\nu(\pi) - \nu'(\pi)) \pi = 0$, with $\nu(\pi) - \nu'(\pi) \neq 0$ for some $\pi \in \text{supp } \nu$, since $\nu \neq \nu'$. Then $\text{supp } \nu$ is linearly dependent, a contradiction. \square

Corollary 6 (Distributions of Posteriors About ω with Affine and Linear Independent Supports).

- i. If $\tau \in \Delta(\Delta(\Omega))$ is Bayes-plausible and $\text{supp } \tau$ is finite, there is a Bayes-plausible $\tau' \in \Delta(\Delta(\Omega))$ such that the posteriors in $\text{supp } \tau'$ are affinely independent and $\text{supp } \tau' \subseteq \text{supp } \tau$.
- ii. If $\{\mu_j\}_{j=1}^J \subset \Delta(\Omega)$ is affinely independent, it is linearly independent.
- iii. If $\tau \in \Delta(\Delta(\Omega))$ is Bayes-plausible and $\text{supp } \tau$ is linearly independent, then $\text{supp } \tau' \neq \text{supp } \tau$ for any Bayes-plausible $\tau' \in \Delta(\Delta(\Omega))$, $\tau' \neq \tau$.

Proof. Follows identically to Lemma 8. \square

Proof of Corollary 1 (Characterization of Optimal Data Publication Mechanisms) From Proposition 1, the maximized value of the designer's problem (2) is $\max_{\nu \in \Delta(K(\epsilon, \pi_0))} \{E_\nu v(\pi) \text{ s.t. } E_\nu \pi = \pi_0\}$. Since $v(\pi) = v_{K(\epsilon, \pi_0)}$ on $K(\epsilon, \pi_0)$, this is equal to $\max_{\nu \in \Delta(K(\epsilon, \pi_0))} \{E_\nu v_{K(\epsilon, \pi_0)}(\pi) \text{ s.t. } E_\nu \pi = \pi_0\}$. By Lemma 6 (ii), $\pi_0 \in \text{ri}(K(\epsilon, \pi_0))$. Since $K(\epsilon, \pi_0) \subseteq \Delta(\Theta)$, $K(\epsilon, \pi_0)$ is bounded in $\mathbb{R}^{2^{N(T+1)}}$; then by Lemma 6 (iii), it is compact and convex. (i) then follows from Proposition 3 in the Online Appendix to Kamenica and Gentzkow (2011).

For (ii), first choose $\nu \in \arg \max_{\nu \in \Delta(K(\epsilon, \pi_0))} \{E_\nu v(\pi) \text{ s.t. } E_\nu \pi = \pi_0\} = \arg \max_{\nu \in \Delta(K(\epsilon, \pi_0))} \{E_\nu v_{K(\epsilon, \pi_0)}(\pi) \text{ s.t. } E_\nu \pi = \pi_0\}$; such a ν exists by (i). If $\text{supp } \nu \not\subseteq \text{ext}(K(\epsilon, \pi_0))$, then by Lemma 7 (i) we can choose a $\hat{\nu}$ with $\text{supp } \hat{\nu} \subseteq \text{ext}(K(\epsilon, \pi_0))$ which is Blackwell-more informative than ν and hence also solves (9). If $\text{supp } \nu \subseteq \text{ext}(K(\epsilon, \pi_0))$, choose $\hat{\nu} = \nu$. Then $\text{supp } \hat{\nu}$ is finite. Moreover, since v is convex, $E_\nu v_{K(\epsilon, \pi_0)}(\pi) = E_\nu v(\pi) \leq E_{\hat{\nu}} v(\pi) = E_{\hat{\nu}} v_{K(\epsilon, \pi_0)}(\pi)$, and $\hat{\nu} \in \arg \max_{\nu \in \Delta(K(\epsilon, \pi_0))} \{E_\nu v_{K(\epsilon, \pi_0)}(\pi) \text{ s.t. } E_\nu \pi = \pi_0\}$ as well.

Now by Lemma 8 (i), there is a Bayes-plausible ν^* such that $\text{supp } \nu^*$ is an affinely independent set and $\text{supp } \nu^* \subseteq \text{supp } \hat{\nu} \subseteq \text{ext}(K(\epsilon, \pi_0))$. Moreover, by Lemma 3 in Yoder (2022), $V_{K(\epsilon, \pi_0)}(\pi) = v_{K(\epsilon, \pi_0)}(\pi) = v(\pi)$ for each $\pi \in \text{supp } \hat{\nu}$, and in addition, $V_{K(\epsilon, \pi_0)}$ is affine on $\text{conv}(\text{supp } \hat{\nu})$: there exists $x \in \mathbb{R}^{2^{N(T+1)}}$ such that for each $\pi \in \text{conv}(\text{supp } \hat{\nu})$, $V_{K(\epsilon, \pi_0)}(\pi) = V_{K(\epsilon, \pi_0)}(\pi_0) + x \cdot (\pi - \pi_0)$. Then since $\text{supp } \nu^* \subseteq \text{supp } \hat{\nu}$, we have $E_{\nu^*} v(\pi) = E_{\nu^*} V_{K(\epsilon, \pi_0)}(\pi) = V_{K(\epsilon, \pi_0)}(\pi_0) + E_{\nu^*}[x \cdot (\pi - \pi_0)] = V_{K(\epsilon, \pi_0)}(\pi_0)$, and so $\nu^* \in \arg \max_{\nu \in \Delta(K(\epsilon, \pi_0))} \{E_\nu v(\pi) \text{ s.t. } E_\nu \pi = \pi_0\}$.

ν^* is induced by the mechanism $(\text{supp } \nu^*, m^*)$ with $m^*(\pi|\theta) = \pi(\theta) \nu^*(\mu) / \pi_0(\theta)$: By Bayes' rule, the posterior probability placed on ω by a decision maker who observes realization π of $(\text{supp } \nu^*, m^*)$ is given by

$$\frac{m^*(\pi|\theta) \pi_0(\theta)}{\sum_{t \in \Theta} m^*(\pi|t) \pi_0(t)} = \frac{\pi(\theta) \nu^*(\pi)}{\sum_{t \in \Theta} \pi(t) \nu^*(\pi)} = \frac{\pi(\theta)}{\sum_{t \in \Theta} \pi(t)}.$$

Then by Proposition 1, $(\text{supp } \nu^*, m^*)$ solves the designer's problem (2). Since it induces ν^* , (iic) follows from affine independence of $\text{supp } \nu^*$ and Lemma 8 (ii); (iicd) then follows from Lemma 8 (iii). And since $\text{supp } \nu^* \subseteq \text{ext}(K(\epsilon, \pi_0))$, (iib) follows from Lemma 7 (i). \square

Oblivious vs. General Publication Mechanisms

Lemma 9. *The following are equivalent:*

- i. *If the distribution $\tau \in \Delta(\Delta(\Omega))$ of posterior beliefs about the population statistic can be induced by an ϵ -differentially private mechanism, it can be induced by an ϵ -differentially private oblivious mechanism.*
- ii. $K_\Omega(\epsilon, \mu_0) = PK(\epsilon, \pi_0)$.

Proof. ((i) \Rightarrow (ii)) We prove the contrapositive. Suppose that $K_\Omega(\epsilon, \mu_0) \neq PK(\epsilon, \pi_0)$. Since oblivious ϵ -differentially private mechanisms are a subset of all ϵ -differentially private mechanisms, by Proposition 2 and Lemma 5, $K_\Omega(\epsilon, \mu_0) \subseteq PK(\epsilon, \pi_0)$. So there must exist $\mu \in PK(\epsilon, \pi_0)$ with $\mu \notin K_\Omega(\epsilon, \mu_0)$; hence, there must exist $\pi \in K(\epsilon, \pi_0)$ with $P\pi \notin K_\Omega(\epsilon, \mu_0)$. By Lemma 6 (ii), there exists $\delta > 0$ such that $\pi' = \pi_0 - \delta(\pi - \pi_0) \in K(\epsilon, \pi_0)$. Then the distribution ν with

$$\nu(\pi) = \frac{\delta}{1 + \delta}, \quad \nu(\pi') = \frac{1}{1 + \delta}$$

is Bayes-plausible with $\text{supp } \nu \subset K(\epsilon, \pi_0)$, and so by Lemma 5 can be induced by an ϵ -differentially private mechanism, which thus induces the distribution of posteriors about the population statistic $\tau = \nu \circ P^{-1} \in \Delta(\Delta(\Omega))$. Then $P\pi \in \text{supp } \tau$, but $P\pi \notin K_\Omega(\epsilon, \mu_0)$; it follows from Proposition 2 that τ cannot be induced with an oblivious ϵ -differentially private mechanism.

((ii) \Rightarrow (i)) Suppose that $K_\Omega(\epsilon, \mu_0) = PK(\epsilon, \pi_0)$, and that $\tau \in \Delta(\Delta(\Omega))$ is the distribution of posterior beliefs about the population statistic induced by the ϵ -differentially private mechanism (S, m) . Then $\text{supp } \tau = P \text{supp } \nu$, where ν is the distribution of posterior beliefs about θ induced by (S, m) . Since (S, m) is ϵ -differentially private, by Lemma 5, $\text{supp } \nu \subset K(\epsilon, \pi_0)$. Then $\text{supp } \tau = P \text{supp } \nu \subset PK(\epsilon, \pi_0) = K_\Omega(\epsilon, \mu_0)$, and by Proposition 2, τ can be induced by an oblivious ϵ -differentially private mechanism. \square

Proof of Theorem 1 (Oblivious Mechanisms and Magnitude Data) Let $d(\theta, \theta') = |\{i | \theta_i \neq \theta'_i\}|$ be the Hamming distance on Θ . For each $\hat{\theta} \in \Theta$, let

$$\pi_{\hat{\theta}}(\theta) = \frac{\pi_0(\theta)e^{-\psi(\theta)\epsilon}}{\sum_{t \in \Theta} \pi_0(t)e^{-\psi(t)\epsilon}}, \text{ where } \psi(\theta) = \begin{cases} \min\{\min\{N, T\}, d(\theta, \hat{\theta})\}, & \omega_\theta \neq \omega_{\hat{\theta}}; \\ \min\{\min\{N, T\} - 1, d(\theta, \hat{\theta})\}, & \omega_\theta = \omega_{\hat{\theta}}. \end{cases}$$

For any θ, θ' with $\theta_{-i} = \theta'_{-i}$ for some i , we have

$$d(\theta, \hat{\theta}) - d(\theta', \hat{\theta}) = \begin{cases} 1, & \theta'_i = \hat{\theta}_i; \\ -1, & \theta_i = \hat{\theta}_i; \\ 0, & \text{otherwise.} \end{cases}$$

and consequently,

$$\psi(\theta) - \psi(\theta') = \begin{cases} 1, & \theta'_i = \hat{\theta}_i \text{ and } d(\theta, \hat{\theta}) \leq \min\{N, T\}; \\ -1, & \theta_i = \hat{\theta}_i \text{ and } d(\theta', \hat{\theta}) \leq \min\{N, T\}; \\ 1, & d(\theta, \hat{\theta}) > \min\{N, T\} \text{ and } \omega_{\theta'} = \omega_{\hat{\theta}}; \\ -1, & d(\theta', \hat{\theta}) > \min\{N, T\} \text{ and } \omega_\theta = \omega_{\hat{\theta}}; \\ 0, & \text{otherwise.} \end{cases}$$

Hence,

$$\left| \log \left(\frac{\pi_{\hat{\theta}}(\theta)}{\pi_{\hat{\theta}}(\theta')} \right) - \log \left(\frac{\pi_0(\theta)}{\pi_0(\theta')} \right) \right| = \left| \log \left(\frac{e^{-\psi(\theta)\epsilon}}{e^{-\psi(\theta')\epsilon}} \right) \right| = |\psi(\theta) - \psi(\theta')|\epsilon \leq \epsilon,$$

and so $\pi_{\hat{\theta}} \in K(\epsilon, \pi_0)$.

Now let $\hat{\theta}$ be a database with $\hat{\theta}_i = 1$ for $i \leq \min\{T, N\}$ and $\hat{\theta}_i = 0$ for $i > \min\{T, N\}$. Then $\omega_{\hat{\theta}} =$

$\min\{T, N\}$. Then (since $\mu_0(0) = \pi_0(\mathbf{0})$)

$$\begin{aligned} \left| \log \left(\frac{P\pi_{\hat{\theta}}(\omega_{\hat{\theta}})}{P\pi_{\hat{\theta}}(0)} \right) - \log \left(\frac{\mu_0(\omega_{\hat{\theta}})}{\mu_0(0)} \right) \right| &= \left| \log \left(\frac{\sum_{\theta: \omega_{\theta} = \omega_{\hat{\theta}}} \pi_0(\theta) e^{-\psi(\theta)\epsilon}}{\mu_0(0) e^{-\min\{T, N\}\epsilon}} \right) - \log \left(\frac{\mu_0(\omega_{\hat{\theta}})}{\mu_0(0)} \right) \right| \\ &= \left| \log \left(\frac{\sum_{\theta: \omega_{\theta} = \omega_{\hat{\theta}}} \pi_0(\theta) e^{-\psi(\theta)\epsilon}}{\mu_0(\omega_{\hat{\theta}}) e^{-\min\{T, N\}\epsilon}} \right) \right|. \end{aligned}$$

Now note that for any θ with $\omega_{\theta} = \omega_{\hat{\theta}} = \min\{T, N\}$, $\psi(\theta) \leq \min\{T, N\} - 1$, and since $T > 1$, $\psi(\hat{\theta}) = 0 < \min\{T, N\} - 1$. Then

$$\log \left(\frac{\sum_{\theta: \omega_{\theta} = \omega_{\hat{\theta}}} \pi_0(\theta) e^{-\psi(\theta)\epsilon}}{\mu_0(\omega_{\hat{\theta}}) e^{-\min\{T, N\}\epsilon}} \right) > \log \left(\frac{e^{\epsilon} \sum_{\theta: \omega_{\theta} = \min\{T, N\}} \pi_0(\theta) e^{-\min\{T, N\}\epsilon}}{\mu_0(\min\{T, N\}) e^{-\min\{T, N\}\epsilon}} \right) = \epsilon.$$

Then $P\pi$ does not satisfy (7) for $\omega = t = \min\{N, T\} \leq T$, and so by definition, $P\pi \notin K_{\Omega}(\epsilon, \mu_0)$. Then $PK(\epsilon, \pi_0) \neq K_{\Omega}(\epsilon, \mu_0)$; the rest of the claim follows from Lemma 9. \square

Proof of Theorem 2 (Oblivious Mechanisms and Binary Data) Suppose $\pi \in K(\epsilon, \pi_0)$. Then for each $\omega \in \Omega \setminus \{0\}$, we have

$$\begin{aligned} P\pi(\omega - 1) &= \sum_{\theta: \omega_{\theta} = \omega - 1} \pi(\theta) = \sum_{\theta: \omega_{\theta} = \omega - 1} \frac{1}{N - (\omega - 1)} \sum_{n: \theta_n = 0} \pi(\theta) \\ &= \frac{1}{(N - (\omega - 1))} \sum_{n=1}^N \sum_{\theta: \theta_n = 0, \omega_{\theta} = \omega - 1} \pi(\theta) = \frac{1}{(N - (\omega - 1))} \sum_{n=1}^N \sum_{\theta': \theta'_n = 1, \omega_{\theta'} = \omega} \pi((0, \theta'_{-n})) \\ &= \frac{1}{(N - (\omega - 1))} \sum_{\theta': \omega_{\theta'} = \omega} \sum_{n: \theta'_n = 1} \pi((0, \theta'_{-n})). \end{aligned}$$

Since respondents are anonymous, if $\omega_{\theta} = \omega_{\theta'}$, then θ' is a permutation of θ , and so, since π_0 is symmetric, $\pi_0(\theta) = \pi_0(\theta')$. It follows that for each θ , $\pi_0(\theta) = \mu_0(\omega_{\theta}) / \binom{N}{\omega_{\theta}}$.

Then since $\pi \in K(\epsilon, \pi_0)$, for each $\theta \in \{0, 1\}^N$, n with $\theta_n = 1$, and $s \in S$, we have

$$\begin{aligned} e^{-\epsilon} \pi(\theta) \frac{\pi_0((0, \theta_{-n}))}{\pi_0(\theta)} &\leq \pi((0, \theta_{-n})) \leq e^{\epsilon} \pi(\theta) \frac{\pi_0((0, \theta_{-n}))}{\pi_0(\theta)} \\ e^{-\epsilon} \pi(\theta) \frac{\mu_0(\omega_{\theta} - 1) \binom{N}{\omega_{\theta}}}{\mu_0(\omega_{\theta}) \binom{N}{\omega_{\theta} - 1}} &\leq \pi((0, \theta_{-n})) \leq e^{\epsilon} \pi(\theta) \frac{\mu_0(\omega_{\theta} - 1) \binom{N}{\omega_{\theta}}}{\mu_0(\omega_{\theta}) \binom{N}{\omega_{\theta} - 1}} \\ e^{-\epsilon} \pi(\theta) \frac{\mu_0(\omega_{\theta} - 1)(N - (\omega_{\theta} - 1))}{\mu_0(\omega_{\theta}) \omega_{\theta}} &\leq \pi((0, \theta_{-n})) \leq e^{\epsilon} \pi(\theta) \frac{\mu_0(\omega_{\theta} - 1)(N - (\omega_{\theta} - 1))}{\mu_0(\omega_{\theta}) \omega_{\theta}} \end{aligned}$$

Hence, for each $\omega \in \Omega \setminus \{0\}$, we have

$$\begin{aligned} e^{-\epsilon} \sum_{\theta': \omega_{\theta'} = \omega} \frac{1}{\omega} \sum_{n: \theta'_n = 1} \pi(\theta') \frac{\mu_0(\omega - 1)}{\mu_0(\omega)} &\leq P\pi(\omega - 1) \leq e^{\epsilon} \sum_{\theta': \omega_{\theta'} = \omega} \frac{1}{\omega} \sum_{n: \theta'_n = 1} \pi(\theta') \frac{\mu_0(\omega - 1)}{\mu_0(\omega)} \\ e^{-\epsilon} \sum_{\theta': \omega_{\theta'} = \omega} \pi(\theta') \frac{\mu_0(\omega - 1)}{\mu_0(\omega)} &\leq P\pi(\omega - 1) \leq e^{\epsilon} \sum_{\theta': \omega_{\theta'} = \omega} \pi(\theta') \frac{\mu_0(\omega - 1)}{\mu_0(\omega)} \\ e^{-\epsilon} P\pi(\omega) \frac{\mu_0(\omega - 1)}{\mu_0(\omega)} &\leq P\pi(\omega - 1) \leq e^{\epsilon} P\pi(\omega) \frac{\mu_0(\omega - 1)}{\mu_0(\omega)}, \end{aligned}$$

and so $P\pi \in K_{\Omega}(\epsilon, \mu_0)$.

Hence, $PK(\epsilon, \pi_0) \subseteq K_{\Omega}(\epsilon, \mu_0)$. And since oblivious ϵ -differentially private mechanisms are a subset of all

ϵ -differentially private mechanisms, by Proposition 2 and Lemma 5, $K_\Omega(\epsilon, \mu_0) \subseteq PK(\epsilon, \pi_0)$. So $PK(\epsilon, \pi_0) = K_\Omega(\epsilon, \mu_0)$; the statement follows by Lemma 9. \square

Proof of Corollary 2 (Differentially Private Data Publication as Information Design) Follows immediately from (6), Proposition 2, and Theorem 2. \square

Proof of Proposition 3 (Oblivious Mechanisms with Two Respondents) Suppose $\pi \in K(\epsilon, \pi_0)$. Then we have

$$\begin{aligned} \frac{P\pi(2)}{\mu_0(2)} &= \frac{\pi((1,1))}{\pi_0((1,1))} \\ e^{-\epsilon}\pi((0,1)) &\leq \pi((1,1)) \frac{\pi_0((0,1))}{\pi_0((1,1))} \leq e^\epsilon\pi((0,1)) \\ e^{-\epsilon}\pi((1,0)) &\leq \pi((1,1)) \frac{\pi_0((1,0))}{\pi_0((1,1))} \leq e^\epsilon\pi((1,0)) \\ \Rightarrow e^{-\epsilon}(\pi((1,0)) + \pi((0,1))) &\leq \pi((1,1)) \frac{\pi_0((1,0)) + \pi_0((0,1))}{\pi_0((1,1))} \leq e^\epsilon(\pi((1,0)) + \pi((0,1))) \\ \Rightarrow e^{-\epsilon}P\pi(1) &\leq P\pi(2) \frac{\pi_0(1)}{\pi_0(2)} \leq e^\epsilon P\pi(1). \end{aligned}$$

Symmetrically, $\frac{P\pi(0)/P\pi(1)}{\pi_0(0)/\pi_0(1)} \in [e^{-\epsilon}, e^\epsilon]$. It follows that $P\pi \in K_\Omega(\epsilon, \mu_0)$.

Then $PK(\epsilon, \pi_0) \subseteq K_\Omega(\epsilon, \mu_0)$. Since oblivious ϵ -differentially private mechanisms are a subset of all ϵ -differentially private mechanisms, by Proposition 2 and Lemma 5, $K_\Omega(\epsilon, \mu_0) \subseteq PK(\epsilon, \pi_0)$. Thus, $PK(\epsilon, \pi_0) = K_\Omega(\epsilon, \mu_0)$; the statement follows by Lemma 9. \square

Proof of Corollary 3 (Characterization of Optimal Oblivious Mechanisms) Follows identically to Corollary 1, relying on Corollary 2, Proposition 2, and Corollaries 4, 5, and 6 instead of Proposition 1, Lemma 5, and Lemmas 6, 7, and 8, respectively.

Geometric Mechanisms and UPRR Dominance

Proof of Lemma 1 (Posteriors Produced by the Geometric Mechanism) Suppose that μ is the decision maker's posterior belief after observing the realization s of the ϵ -geometric data publication mechanism. By Bayes' rule, for each $\omega \in \Omega$,

$$\mu(\omega) = \frac{\sigma_\epsilon^g(s|\omega)\mu_0(\omega)}{\sum_{\omega' \in \Omega} \sigma_\epsilon^g(s|\omega')\mu_0(\omega')}.$$

Then for each $\omega \in \Omega \setminus \{0\}$, we have

$$\log \left(\frac{\mu(\omega)/\mu(\omega-1)}{\mu_0(\omega)/\mu_0(\omega-1)} \right) = \log \left(\frac{\sigma_\epsilon^g(s|\omega)}{\sigma_\epsilon^g(s|\omega-1)} \right) = -\epsilon|s-\omega| + \epsilon|s-(\omega-1)| = \begin{cases} \epsilon, & s \geq \omega; \\ -\epsilon, & s < \omega. \end{cases}$$

The claim follows. \square

Proof of Lemma 2 (Transitivity of UPRR Dominance) Fix $\mu \in \text{supp } \tau$, and let $\omega^*(\mu)$ be its peak used to show $\tau \succeq_{\text{UPRR}} \tau'$.

We first show that there exists $\mu' \in \text{supp } \tau'$ and $\omega' \leq \omega^*(\mu)$ such that for all $\mu'' \in \text{supp } \tau''$, $\mu'(\omega)/\mu''(\omega)$ is nonincreasing on $[\omega', \infty)$. Suppose not. For each $\mu' \in \text{supp } \tau'$, let $\omega^*(\mu')$ be the peak used to show $\tau' \succeq_{\text{UPRR}} \tau''$. Then for each $\mu'' \in \text{supp } \tau''$, $\mu'(\omega)/\mu''(\omega)$ is nondecreasing on $(-\infty, \omega^*(\mu)]$ and nonincreasing on $[\omega^*(\mu), \infty)$. Then we must have $\omega^*(\mu') > \omega^*(\mu)$, or we could choose μ' and $\omega' = \omega^*(\mu')$ and arrive at a contradiction. Then $\frac{\mu'(\omega)}{\mu''(\omega)} \geq \frac{\mu'(\omega^*(\mu))}{\mu''(\omega^*(\mu))}$ for each $\mu' \in \text{supp } \tau'$, each $\omega \in [\omega^*(\mu), \omega^*(\mu')]$, and each $\mu'' \in$

$\text{supp } \tau''$. Moreover, for each $\mu' \in \text{supp } \tau'$, this inequality must be strict for some ω (in particular, $\omega = \omega^*(\mu')$) and some $\mu'' \in \text{supp } \tau''$, or we could choose μ' and $\omega' = \omega^*(\mu)$ and arrive at a contradiction.

Let $\hat{\mu}' = \arg \min_{\mu' \in \text{supp } \tau'} \omega^*(\mu')$. Then for all $\mu' \in \text{supp } \tau'$ and $\mu'' \in \text{supp } \tau''$, we have $\frac{\mu'(\omega^*(\hat{\mu}'))}{\mu''(\omega^*(\hat{\mu}'))} \geq \frac{\mu'(\omega^*(\mu))}{\mu''(\omega^*(\mu))}$, and for some $\mu'' \in \text{supp } \tau''$, we have

$$\begin{aligned} \frac{\hat{\mu}'(\omega^*(\hat{\mu}'))}{\mu''(\omega^*(\hat{\mu}'))} &> \frac{\hat{\mu}'(\omega^*(\mu))}{\mu''(\omega^*(\mu))} \\ \Rightarrow \frac{\mu_0(\omega^*(\hat{\mu}'))}{\mu''(\omega^*(\hat{\mu}'))} &= \frac{\sum_{\mu' \in \text{supp } \tau'} \tau'(\mu') \mu'(\omega^*(\hat{\mu}'))}{\mu''(\omega^*(\hat{\mu}'))} > \frac{\sum_{\mu' \in \text{supp } \tau'} \tau'(\mu') \mu'(\omega^*(\mu))}{\mu''(\omega^*(\mu))} = \frac{\mu_0(\omega^*(\mu))}{\mu''(\omega^*(\mu))} \end{aligned}$$

Likewise, we have $\frac{\mu_0(\omega^*(\hat{\mu}'))}{\mu''(\omega^*(\hat{\mu}'))} \geq \frac{\mu_0(\omega^*(\mu))}{\mu''(\omega^*(\mu))}$ for all $\mu'' \in \text{supp } \tau''$. Taking reciprocals yields

$$1 = \frac{\mu_0(\omega^*(\hat{\mu}'))}{\mu_0(\omega^*(\mu))} = \frac{\sum_{\mu'' \in \text{supp } \tau''} \tau''(\mu'') \mu''(\omega^*(\hat{\mu}'))}{\mu_0(\omega^*(\hat{\mu}'))} > \frac{\sum_{\mu'' \in \text{supp } \tau''} \tau''(\mu'') \mu''(\omega^*(\mu))}{\mu_0(\omega^*(\mu))} = \frac{\mu_0(\omega^*(\mu))}{\mu_0(\omega^*(\mu))} = 1,$$

a contradiction.

So there exists $\mu' \in \text{supp } \tau'$ and $\omega' \leq \omega^*(\mu)$ such that for all $\mu'' \in \text{supp } \tau''$ and $\hat{\omega} \geq \omega \geq \omega'$, $\mu'(\hat{\omega})/\mu''(\hat{\omega}) \leq \mu'(\omega)/\mu''(\omega)$. Since $\tau \succeq_{\text{UPRR}} \tau'$, $\mu(\hat{\omega})/\mu'(\hat{\omega}) \leq \mu(\omega)/\mu'(\omega)$. Since all the terms are nonnegative, it follows that $\mu(\hat{\omega})/\mu''(\hat{\omega}) \leq \mu(\omega)/\mu''(\omega)$ for each $\hat{\omega} \geq \omega \geq \omega^*(\mu)$ and $\mu'' \in \text{supp } \tau''$.

A symmetric argument shows that $\mu(\hat{\omega})/\mu''(\hat{\omega}) \geq \mu(\omega)/\mu''(\omega)$ for each $\omega \leq \hat{\omega} \leq \omega^*(\mu)$ and $\mu'' \in \text{supp } \tau''$. Since these statements hold for each $\mu \in \text{supp } \tau$, it follows that $\tau \succeq_{\text{UPRR}} \tau''$. \square

Proof of Lemma 3 (Geometric Mechanisms are UPRR-Dominant) Suppose that $\mu \in \text{supp } \tau_e^g$. Then by Lemma 1, for each $\omega \in (0, \omega^*(\mu)]$, $\log \left(\frac{\mu(\omega)/\mu_0(\omega)}{\mu(\omega-1)/\mu_0(\omega-1)} \right) = \epsilon$, or equivalently, $\mu(\omega)/\mu(\omega-1) = e^\epsilon \mu_0(\omega)/\mu_0(\omega-1)$. Hence, for each $\omega' \leq \omega'' \leq \omega^*(\mu)$,

$$\mu(\omega'')/\mu(\omega') = \prod_{i=\omega'+1}^{\omega''} \mu(\omega_i)/\mu(\omega_i-1) = e^{\epsilon(\omega''-\omega')} \mu_0(\omega'')/\mu_0(\omega').$$

Likewise, for each $\omega > \omega^*(\mu)$, we have $\log \left(\frac{\mu(\omega)/\mu_0(\omega)}{\mu(\omega-1)/\mu_0(\omega-1)} \right) = -\epsilon$, or equivalently, $\mu(\omega)/\mu(\omega-1) = e^{-\epsilon} \mu_0(\omega)/\mu_0(\omega-1)$. Hence, for each $\omega'' \geq \omega' \geq \omega^*(\mu)$,

$$\mu(\omega'')/\mu(\omega') = \prod_{i=\omega'+1}^{\omega''} \mu(\omega_i)/\mu(\omega_i-1) = e^{-\epsilon(\omega''-\omega')} \mu_0(\omega'')/\mu_0(\omega').$$

Fix $\mu' \in \text{supp } \tau$. Since τ is induced by an oblivious ϵ -differentially private mechanism, by Proposition 2, $\mu' \in K_\Omega(\epsilon, \mu_0)$. Then for each $\omega \in \Omega \setminus \{0\}$, $\left| \log \left(\frac{\mu'(\omega)/\mu_0(\omega)}{\mu'(\omega-1)/\mu_0(\omega-1)} \right) \right| \leq \epsilon$, or equivalently, $e^{-\epsilon} \mu_0(\omega)/\mu_0(\omega-1) \leq \mu'(\omega)/\mu'(\omega-1) \leq e^\epsilon \mu_0(\omega)/\mu_0(\omega-1)$. Hence, for each $\omega' \leq \omega''$,

$$\mu'(\omega'')/\mu'(\omega') = \prod_{i=\omega'+1}^{\omega''} \mu'(\omega_i)/\mu'(\omega_i-1) \in \left[e^{-\epsilon(\omega''-\omega')} \frac{\mu_0(\omega'')}{\mu_0(\omega')}, e^{\epsilon(\omega''-\omega')} \frac{\mu_0(\omega'')}{\mu_0(\omega')} \right].$$

Then for each $\omega' \leq \omega'' \leq \omega^*(\mu)$, $\mu(\omega'')/\mu(\omega') = e^{\epsilon(\omega''-\omega')} \mu_0(\omega'')/\mu_0(\omega') \geq \mu'(\omega'')/\mu'(\omega')$; hence, $\mu(\omega'')/\mu'(\omega'') \geq \mu(\omega')/\mu'(\omega')$. Likewise, for each $\omega'' \geq \omega' \geq \omega^*(\mu)$, $\mu(\omega'')/\mu(\omega') = e^{-\epsilon(\omega''-\omega')} \mu_0(\omega'')/\mu_0(\omega') \leq \mu'(\omega'')/\mu'(\omega')$; hence, $\mu(\omega'')/\mu'(\omega'') \leq \mu(\omega')/\mu'(\omega')$. \square

Proof of Lemma 4 (Equivalence of Frechét Representations) (i): The marginal cdf for X , F_X , is given by

$$\begin{aligned} F_X(x) &= F(N, x) = \int_0^x \sum_{y=0}^N \sum_{\mu \in t^{-1}(Q_{\tau,t}(z))} \frac{\tau(\mu)\mu(y)}{r_{\tau,t}(Q_{\tau,t}(z))} dz = \int_0^x \sum_{\mu \in t^{-1}(Q_{\tau,t}(z))} \frac{\tau(\mu)}{r_{\tau,t}(Q_{\tau,t}(z))} \sum_{y=0}^N \mu(y) dz \\ &= \int_0^x \sum_{\mu \in t^{-1}(Q_{\tau,t}(z))} \frac{\tau(\mu)}{\tau(t^{-1}(Q_{\tau,t}(z)))} dz \text{ (by definition of } r_{\tau,t}) \\ &= \int_0^x 1 dz = x, \end{aligned}$$

and so $X \sim U([0, 1])$. The marginal cdf for W , F_W , is given by

$$\begin{aligned} F_W(w) &= F(w, 1) = \int_0^1 \sum_{y=0}^w \sum_{\mu \in t^{-1}(Q_{\tau,t}(z))} \frac{\tau(\mu)\mu(y)}{r_{\tau,t}(Q_{\tau,t}(z))} dz = \sum_{y=0}^w \int_0^1 \sum_{\mu \in t^{-1}(Q_{\tau,t}(z))} \frac{\tau(\mu)\mu(y)}{r_{\tau,t}(Q_{\tau,t}(z))} dz \\ &= \sum_{y=0}^w E \left[\sum_{\mu \in t^{-1}(Q_{\tau,t}(X))} \frac{\tau(\mu)\mu(y)}{r_{\tau,t}(Q_{\tau,t}(X))} \right] \text{ (since } X \sim U([0, 1]) \text{)}. \end{aligned}$$

Since $X \sim U([0, 1])$ and $Q_{\tau,t}$ is an inverse cdf of T , $Q_{\tau,t}(X)$ is equal in distribution to T . Then

$$\begin{aligned} F_W(w) &= \sum_{y=0}^w \sum_{z \in t(\text{supp } \tau)} \sum_{\mu \in t^{-1}(z)} \frac{\tau(\mu)\mu(y)}{r_{\tau,t}(z)} r_{\tau,t}(z) = \sum_{y=0}^w \sum_{z \in t(\text{supp } \tau)} \sum_{\mu \in t^{-1}(z)} \tau(\mu)\mu(y) \\ &= \sum_{y=0}^w \sum_{\mu \in \text{supp } \tau} \tau(\mu)\mu(y) = \sum_{y=0}^w \mu_0(y) \text{ (since } \tau \text{ is Bayes-plausible)} \end{aligned}$$

and so $W \sim \mu_0$.

(ii): By the law of iterated expectations, we have

$$E[h(W, Q_{\tau,t}(X))] = E[E[h(W, Q_{\tau,t}(X))|X]] = E \left[\sum_{w \in \Omega} \sum_{\mu \in t^{-1}(Q_{\tau,t}(X))} \frac{\tau(\mu)\mu(w)}{r_{\tau,t}(Q_{\tau,t}(X))} h(w, Q_{\tau,t}(X)) \right].$$

By (i), $X \sim U([0, 1])$. Then since $Q_{\tau,t}$ is an inverse cdf, $Q_{\tau,t}(X) \sim R_{\tau,t}$, and we have

$$\begin{aligned} E[h(W, Q_{\tau,t}(X))] &= \sum_{z \in t(\text{supp } \tau)} r_{\tau,t}(z) \sum_{w \in \Omega} \sum_{\mu \in t^{-1}(z)} \frac{\tau(\mu)\mu(w)}{r_{\tau,t}(z)} h(w, z) \\ &= \sum_{w \in \Omega} \sum_{z \in t(\text{supp } \tau)} \sum_{\mu \in t^{-1}(z)} \tau(\mu)\mu(w) h(w, z) = \sum_{w \in \Omega} \sum_{z \in t(\text{supp } \tau)} \sum_{\mu \in t^{-1}(z)} \tau(\mu)\mu(w) h(w, t(\mu)) \\ &= \sum_{w \in \Omega} \sum_{\mu \in \text{supp } \tau} \tau(\mu)\mu(w) h(w, t(\mu)) = E_{\tau}[E_{\mu}h(w, t(\mu))], \end{aligned}$$

as desired.

(iii): Since A is compact and h is upper semicontinuous, $\max_{a \in A} (\sum_{w \in \Omega} \mu(w) h(w, a)) = \max_{a \in A} E_{\mu}[h(w, a)]$

exists for each $\mu \in \Delta(\Omega)$. By the law of iterated expectations, we have

$$\begin{aligned}
E[h(W, Q_{\tau',s}(X))] &= E[E[h(W, Q_{\tau',s}(X))|X]] = E\left[\sum_{w \in \Omega} \sum_{\mu \in t^{-1}(Q_{\tau,t}(X))} \frac{\tau(\mu)\mu(w)}{r_{\tau,t}(Q_{\tau,t}(X))} h(w, Q_{\tau',s}(X))\right] \\
&= E\left[\sum_{\mu \in t^{-1}(Q_{\tau,t}(X))} \frac{\tau(\mu)}{r_{\tau,t}(Q_{\tau,t}(X))} \sum_{w \in \Omega} \mu(w) h(w, Q_{\tau',s}(X))\right] \\
&\leq E\left[\sum_{\mu \in t^{-1}(Q_{\tau,t}(X))} \frac{\tau(\mu)}{r_{\tau,t}(Q_{\tau,t}(X))} \max_{a \in A} \left(\sum_{w \in \Omega} \mu(w) h(w, a)\right)\right] = E\left[\sum_{\mu \in t^{-1}(Q_{\tau,t}(X))} \frac{\tau(\mu)}{r_{\tau,t}(Q_{\tau,t}(X))} \max_{a \in A} E_{\mu}[h(\omega, a)]\right].
\end{aligned}$$

By (i), $X \sim U([0, 1])$. Then since $Q_{\tau,t}$ is an inverse cdf, $Q_{\tau,t}(X) \sim R_{\tau,t}$, and we have

$$\begin{aligned}
E[h(W, Q_{\tau',s}(X))] &\leq \sum_{z \in t(\text{supp } \tau)} r_{\tau,t}(z) \sum_{\mu \in t^{-1}(z)} \frac{\tau(\mu)}{r_{\tau,t}(z)} \max_{a \in A} E_{\mu}[h(\omega, a)] \\
&= \sum_{z \in t(\text{supp } \tau)} \sum_{\mu \in t^{-1}(z)} \tau(\mu) \max_{a \in A} E_{\mu}[h(\omega, a)] \\
&= \sum_{\mu \in \text{supp } \tau} \tau(\mu) \max_{a \in A} E_{\mu}[h(\omega, a)] = E_{\tau}[\max_{a \in A} E_{\mu}[h(\omega, a)]],
\end{aligned}$$

as desired. \square

Proof of Proposition 5 (UPRR-Dominance Implies Supermodular Stochastic Dominance) Let F and G denote the cumulative distribution functions of (W, X) and (\hat{W}, \hat{X}) , respectively. Further, let

$$\begin{aligned}
f(w, x) &= \sum_{\mu \in \omega^{*-1}(Q_{\tau,\omega^*}(x))} \frac{\tau(\mu)\mu(w)}{q_{\tau,\omega^*}(Q_{\tau,\omega^*}(x))}; \\
g(w, x) &= \sum_{\mu \in t^{-1}(Q_{\tau',t}(x))} \frac{\tau'(\mu)\mu(w)}{r_{\tau',t}(Q_{\tau',t}(x))}.
\end{aligned}$$

Since $\tau \succeq_{\text{UPRR}} \tau'$, for any $x \in [0, 1]$, $\omega' \leq \omega'' \leq Q_{\tau,\omega^*}(x) \leq \tilde{\omega} \leq \tilde{\tilde{\omega}}$, $\mu \in \omega^{*-1}(Q_{\tau,\omega^*}(x))$, and $\mu' \in \text{supp } \tau'$, we have $\frac{\mu(\omega'')}{\mu(\omega')} \geq \frac{\mu'(\omega'')}{\mu'(\omega')}$ and $\frac{\mu(\tilde{\tilde{\omega}})}{\mu(\tilde{\omega})} \leq \frac{\mu'(\tilde{\tilde{\omega}})}{\mu'(\tilde{\omega})}$, or equivalently,

$$\mu(\omega'')\mu'(\omega') \geq \mu'(\omega'')\mu(\omega') \quad \text{and} \quad \mu(\tilde{\tilde{\omega}})\mu'(\tilde{\omega}) \leq \mu'(\tilde{\tilde{\omega}})\mu(\tilde{\omega}).$$

It follows that for any $x \in [0, 1]$, $\omega' \leq \omega'' \leq Q_{\tau,\omega^*}(x) \leq \tilde{\omega} \leq \tilde{\tilde{\omega}}$, and $\mu' \in \text{supp } \tau'$,

$$\begin{aligned}
f(\omega'', x)\mu'(\omega') - \mu'(\omega'')f(\omega', x) &= \sum_{\mu \in \omega^{*-1}(Q_{\tau,\omega^*}(x))} \frac{\tau(\mu)(\mu(\omega'')\mu'(\omega') - \mu(\omega')\mu'(\omega''))}{q_{\tau,\omega^*}(Q_{\tau,\omega^*}(x))} \geq 0, \\
f(\tilde{\tilde{\omega}}, x)\mu'(\tilde{\omega}) - \mu'(\tilde{\tilde{\omega}})f(\tilde{\omega}, x) &= \sum_{\mu \in \omega^{*-1}(Q_{\tau,\omega^*}(x))} \frac{\tau(\mu)(\mu(\tilde{\tilde{\omega}})\mu'(\tilde{\omega}) - \mu(\tilde{\omega})\mu'(\tilde{\tilde{\omega}}))}{q_{\tau,\omega^*}(Q_{\tau,\omega^*}(x))} \leq 0.
\end{aligned}$$

Since Q_{τ,ω^*} is an inverse cdf, it is nondecreasing; then for each $x \in [0, 1]$, $y \in [x, 1]$, $z \in [0, x]$, and $\omega' \leq \omega'' \leq Q_{\tau,\omega^*}(x) \leq \tilde{\omega} \leq \tilde{\tilde{\omega}}$, we have $\omega' \leq \omega'' \leq Q_{\tau,\omega^*}(y)$ and $Q_{\tau,\omega^*}(z) \leq \tilde{\omega} \leq \tilde{\tilde{\omega}}$, and so for any $\mu' \in \text{supp } \tau'$,

$$\mu(\omega'')\mu'(\omega') \geq \mu'(\omega'')\mu(\omega', y) \quad \text{and} \quad f(\tilde{\tilde{\omega}}, z)\mu'(\tilde{\omega}) \leq \mu'(\tilde{\tilde{\omega}})f(\tilde{\omega}, z),$$

and thus, for each $y', z' \in [0, 1]$,

$$f(\omega'', y)g(\omega', y') - g(\omega'', y')f(\omega', y) = \sum_{\mu' \in t^{-1}(Q_{\tau', t}(y'))} \frac{\tau'(\mu)(f(\omega'', y)\mu'(\omega') - \mu'(\omega'')f(\omega', y))}{r_{\tau', t}(Q_{\tau', t}(y'))} \geq 0, \quad (14)$$

$$f(\tilde{\omega}, z)g(\tilde{\omega}, z') - g(\tilde{\omega}, z')f(\tilde{\omega}, z) = \sum_{\mu' \in t^{-1}(Q_{\tau', t}(z'))} \frac{\tau'(\mu)(f(\tilde{\omega}, z)\mu'(\tilde{\omega}) - \mu'(\tilde{\omega})f(\tilde{\omega}, z))}{r_{\tau', t}(Q_{\tau', t}(z'))} \leq 0. \quad (15)$$

Integrating (14) on $[x, 1] \times [x, 1]$ and (15) on $[0, x] \times [0, x]$ yields

$$\begin{aligned} \int_x^1 f(\omega'', y)dy \int_x^1 g(\omega', y')dy' &\geq \int_x^1 g(\omega'', y')dy' \int_x^1 f(\omega', y)dy \\ &\Leftrightarrow \frac{\int_x^1 f(\omega'', y)dy}{\int_x^1 g(\omega'', y)dy} \geq \frac{\int_x^1 f(\omega', y)dy}{\int_x^1 g(\omega', y)dy}; \end{aligned} \quad (16)$$

$$\begin{aligned} \int_0^x f(\tilde{\omega}, z)dz \int_0^x g(\tilde{\omega}, z')dz' &\leq \int_0^x g(\tilde{\omega}, z')dz' \int_0^x f(\tilde{\omega}, z)dz \\ &\Leftrightarrow \frac{\int_0^x f(\tilde{\omega}, z)dz}{\int_0^x g(\tilde{\omega}, z)dz} \leq \frac{\int_0^x f(\tilde{\omega}, z)dz}{\int_0^x g(\tilde{\omega}, z)dz}. \end{aligned} \quad (17)$$

It follows from (17) that for each $x \in [0, 1]$, $\int_0^x g(\omega, z)dz$ single-crosses $\int_0^x f(\omega, z)dz$ on $[Q_{\tau, \omega^*}(x), N]$: for all $\tilde{\omega} \in [Q_{\tau, \omega^*}(x), N]$ and $\tilde{\omega} \in [\tilde{\omega}, N]$,

$$\int_0^x g(\tilde{\omega}, z)dz \geq \int_0^x f(\tilde{\omega}, z)dz \Rightarrow \int_0^x g(\tilde{\omega}, z)dz \geq \int_0^x f(\tilde{\omega}, z)dz. \quad (18)$$

Likewise, it follows from (16) that for each $x \in [0, 1]$, $\int_x^1 f(\omega, z)dz$ single-crosses $\int_x^1 g(\omega, z)dz$ on $[0, Q_{\tau, \omega^*}(x)]$: for all $\omega' \in [0, Q_{\tau, \omega^*}(x)]$ and $\omega'' \in [\omega', Q_{\tau, \omega^*}(x)]$,

$$\int_x^1 f(\omega', z)dz \geq \int_x^1 g(\omega', z)dz \Rightarrow \int_x^1 f(\omega'', z)dz \geq \int_x^1 g(\omega'', z)dz. \quad (19)$$

By Lemma 4 (i), F and G are both elements of $\mathcal{M}(\mu_0, U([0, 1]))$, and so $\int_0^1 g(\omega, z)dz = \int_0^1 f(\omega, z)dz = \mu_0(\omega)$ for each $\omega \in \Omega$. Hence, for each $\omega \in \Omega$,

$$\begin{aligned} \int_x^1 f(\omega, z)dz &\geq \int_x^1 g(\omega, z)dz \Leftrightarrow \mu_0(\omega) - \int_0^x f(\omega, z)dz \geq \mu_0(\omega) - \int_0^x g(\omega, z)dz \\ &\Leftrightarrow \int_0^x g(\omega, z)dz \geq \int_0^x f(\omega, z)dz. \end{aligned}$$

Then (19) implies that for each $x \in [0, 1]$, $\int_0^x g(\omega, z)dz$ single-crosses $\int_0^x f(\omega, z)dz$ on $[0, Q_{\tau, \omega^*}(x)]$: for all $\omega' \in [0, Q_{\tau, \omega^*}(x)]$ and $\omega'' \in [\omega', Q_{\tau, \omega^*}(x)]$,

$$\int_0^x g(\omega', z)dz \geq \int_0^x f(\omega', z)dz \Rightarrow \int_0^x g(\omega'', z)dz \geq \int_0^x f(\omega'', z)dz. \quad (20)$$

Moreover, for each $x \in [0, 1]$, $\tilde{\omega} \in [0, Q_{\tau, \omega^*}(x)]$, and $\tilde{\omega} \in [Q_{\tau, \omega^*}(x), N]$, (18) and (20) imply that

$$\begin{aligned} \int_0^x g(\tilde{\omega}, z)dz &\geq \int_0^x f(\tilde{\omega}, z)dz \Rightarrow \int_0^x g(Q_{\tau, \omega^*}(x), z)dz \geq \int_0^x f(Q_{\tau, \omega^*}(x), z)dz \\ &\Rightarrow \int_0^x g(\tilde{\omega}, z)dz \geq \int_0^x f(\tilde{\omega}, z)dz. \end{aligned} \quad (21)$$

Then from (18), (20), and (21), for each $x \in [0, 1]$, $\int_0^x g(\omega, z)dz$ single-crosses $\int_0^x f(\omega, z)dz$ on all of Ω : (18) holds for all $\tilde{\omega}, \tilde{\omega} \in \Omega$ with $\tilde{\omega} \leq \tilde{\omega}$. Hence, there exists $\omega_0(x) \in \mathbb{R}$ such that for each $\omega \in \Omega$ with $\omega < \omega_0(x)$,

$\int_0^x f(\omega, z)dz - \int_0^x g(\omega, z)dz \geq 0$, and for each $\omega \in \Omega$ with $\omega \geq \omega_0(x)$, $\int_0^x f(\omega, z)dz - \int_0^x g(\omega, z)dz \leq 0$.

Then for all $w \in \Omega, x \in [0, 1]$ with $w < \omega_0(x)$, we have

$$F(w, x) - G(w, x) = \sum_{y=0}^w \left(\int_0^x f(y, z)dz - \int_0^x g(y, z)dz \right) \geq 0.$$

Moreover, since $X \sim U([0, 1])$ and $\hat{X} \sim U([0, 1])$ by Lemma 4 (i), $F(N, x) - G(N, x) = F_X(x) - G_{\hat{X}}(x) = 0$. Then for all $w \in \Omega, x \in [0, 1]$ with $\omega_0(x) \leq w < N$, we have

$$\begin{aligned} F(w, x) - G(w, x) &= F(N, x) - G(N, x) - \sum_{y=w+1}^N \left(\int_0^x f(y, z)dz - \int_0^x g(y, z)dz \right) \\ &= - \sum_{y=w+1}^N \left(\int_0^x f(y, z)dz - \int_0^x g(y, z)dz \right) \geq 0. \end{aligned}$$

Hence, $F(w, x) \geq G(w, x)$ for all $w \in \Omega$ and $x \in [0, 1]$. Equivalently (see Shaked and Shanthikumar (2007) (9.A.18); Epstein and Tanny (1980) Theorem 6) for every supermodular $h : \Omega \times [0, 1] \rightarrow \mathbb{R}$,

$$E_F[h(w, x)] \geq E_G[h(w, x)].$$

Proof of Theorem 4 (UPRR-Dominance Implies Dominance in Supermodular Problems) Since h is upper semicontinuous, so is $\sum_{w \in \Omega} \mu(w)h(w, a)$ for any $\mu \in \Delta(\Omega)$. Then since A is compact, $\arg \max_{a \in A} E_\mu[h(\omega, a)] = \arg \max_{a \in A} \{\sum_{w \in \Omega} \mu(w)h(w, a)\}$ is nonempty.

For each $\mu \in \text{supp } \tau'$, choose $a^*(\mu) \in \arg \max_{a \in A} E_\mu[h(\omega, a)]$, and let (\hat{W}, \hat{X}) be the Frechét representation of τ' with respect to a^* . Then by Lemma 4 (ii),

$$E[h(\hat{W}, Q_{\tau', a^*}(\hat{X}))] = E_{\tau'} [E_\mu[h(\omega, a^*(\mu))]] = E_{\tau'} \left[\max_{a \in A} E_\mu[h(\omega, a)] \right]. \quad (22)$$

Since it is an inverse cdf, Q_{τ', a^*} is nondecreasing, and so for any $x', x'' \in [0, 1]$, $Q_{\tau', a^*}(x' \wedge x'') = Q_{\tau', a^*}(x') \wedge Q_{\tau', a^*}(x'')$ and $Q_{\tau', a^*}(x' \vee x'') = Q_{\tau', a^*}(x') \vee Q_{\tau', a^*}(x'')$. It follows that $h(w, Q_{\tau', a^*}(x))$ is supermodular: For any $x', x'' \in [0, 1]$ and $w', w'' \in \Omega$, supermodularity of h implies

$$\begin{aligned} h(w', Q_{\tau', a^*}(x')) + h(w'', Q_{\tau', a^*}(x'')) \\ \geq h(w'' \vee w', Q_{\tau', a^*}(x') \vee Q_{\tau', a^*}(x'')) + h(w'' \wedge w', Q_{\tau', a^*}(x') \wedge Q_{\tau', a^*}(x'')) \\ = h(w'' \vee w', Q_{\tau', a^*}(x'' \vee x')) + h(w'' \wedge w', Q_{\tau', a^*}(x'' \wedge x')). \end{aligned}$$

Let $\omega^* : \text{supp } \tau \rightarrow \mathbb{R}$ be the function mapping each posterior in $\text{supp } \tau$ to its peak used in the UPRR ordering $\tau \succeq_{\text{UPRR}} \tau'$, and let (W, X) be the Frechét representation of τ with respect to ω^* . Then by Proposition 5 and (22),

$$E[h(W, Q_{\tau', a^*}(X))] \geq E[h(\hat{W}, Q_{\tau', a^*}(\hat{X}))] = E_{\tau'} \left[\max_{a \in A} E_\mu[h(\omega, a)] \right].$$

The statement then follows directly from Lemma 4 (iii). \square

Proof of Theorem 3 (Optimality of the Geometric Mechanism for Supermodular Problems) By Corollary 3 (iic), the designer's problem (2) is solved by an ϵ -differentially private oblivious mechanism which induces a distribution of posteriors about the population statistic τ^* whose support is linearly independent, and hence finite. By Lemma 3, $\tau_\epsilon^\delta \succeq_{\text{UPRR}} \tau^*$. The claim then follows from Theorem 4. \square