Information Design for Differential Privacy

Ian M. Schmutte and Nathan Yoder*

December 6, 2022

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 this is a constrained information design problem, and characterize its solution. When the underlying database is drawn from a symmetric distribution — for instance, if individuals' data are i.i.d. — we show that the problem's dimensionality can be reduced, and that its solution belongs to a simpler class of mechanisms. 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, Department of Economics and U.S. Census Bureau; E-mail: schmutte@uga.edu. Yoder: University of Georgia, Terry College of Business, 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, R. Vijay Krishna, Meg Meyer, Marzena Rostek, Nikhil Vellodi, 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, and the 2022 ASU Theory Conference 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. The views in this paper are those of the authors and not necessarily those of the U.S. Census Bureau.

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). Statistical agencies and other data providers must find ways to release information about data that are as informative as possible about important population characteristics but that also protect the privacy of individuals.

To address these issues, technology firms like Google, LinkedIn, and Uber, and statistical agencies like the U.S. Census Bureau, have begun 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). 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's data can change the probability of producing any given output.

However, there is little clarity about how to maximize the value of statistics to data users while preserving privacy. Our paper sheds light on this question by explicitly modeling the problem faced by a data provider that must publish information while satisfying a differential privacy requirement. The data provider has a database with responses collected from individuals about whether they have a certain trait. Its objective is to maximize the ex ante welfare of data users whose payoffs depend on their action and a population statistic — the number of proportion of respondents that have the trait in question. Importantly, the data provider must decide how to publish before accessing the data. It does so by choosing a *publication mechanism*, or map from data to published output. Because the data provider faces a differential privacy requirement, this map must be stochastic: given a dataset, it generates a probability distribution whose realization is then published.

The data provider's problem in our model mirrors problems faced by 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 differential privacy to protect individuals'

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

[2] Abowd 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 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 prod-

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 valid differential privacy guarantee requires data providers like these to commit to a mechanism without reference to the underlying data: Unless it is formally accounted for as part of the mechanism, using the data to develop or calibrate the mechanism violates the assumptions required to prove it satisfies differential privacy in the first place.[3]

Because the data provider must commit to a signal (its publication mechanism) about an unknown parameter (the database), it faces a problem of the type considered in the extensive literature on *information design* (Kamenica and Gentzkow, 2011; Kamenica, 2019). Our first main result shows that we can use key tools from this literature — i.e., considering a signal as a distribution of posterior beliefs, and concavifying a value function over those beliefs — to characterize the problem and its solution. Specifically, Theorem 1 shows that the data provider faces a standard information design problem, subject to a set of linear constraints on the posterior beliefs that the publication mechanism can induce. Because the set of posteriors that satisfy those constraints is convex, the characterization from Kamenica and Gentzkow (2011) extends naturally to differential privacy. And since the data provider shares the data user's preferences, we can go further: Blackwell's (1953) theorem tells us we can restrict attention to the extreme points of the constraint set, which we show must cause the differential privacy requirement to bind in a specific way.

Since the data user's payoffs only depend on the value of a population statistic, we might think that the data provider can restrict attention to mechanisms that depend only on the value of that statistic, rather than the whole database. This class of mechanisms — which the computer science literature calls *oblivious* — includes several mechanisms commonly used in practice to publish information about the kind of data considered in this paper. Our second main result, Theorem 2, shows that when the respondents are *anonymous*, in the sense that the data user's prior belief about their traits is symmetric, this conjecture is correct: Oblivious mechanisms are without loss. But we also give a counterexample (Example 2) showing that when the individuals in the data are *not* anonymous — for instance, if there is outside information available about their traits — oblivious mechanisms need not be optimal. In such cases, using that outside information to design the mechanism

ucts 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*.

[3]For 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.

can allow data providers to make their publications more useful without compromising privacy.

Our third main result shows that in a large class of applications, we can move beyond the broad characterization of Theorems 1 and 2, and derive a practical mechanism that solves the data provider's problem. Specifically, Theorem 3 shows that the *ϵ-geometric mechanism* (Ghosh et al., 2012) solves the data provider's problem in any setting where data users' actions can be ordered from highest to lowest, their payoff functions are supermodular, and respondents are anonymous. This mechanism is simple to implement: Under the geometric mechanism, the published data is just the sum of the true population statistic and a two-sided geometrically distributed noise term. It is widely used by organizations like the U.S. Census Bureau to publish information about the kinds of population statistics considered in our model.[4]

The decision problems considered in Theorem 3 — those in which higher actions and higher population statistics are complementary — are common in practice. For instance, if an individual believes that there is a greater incidence in her community of active infections of a pandemic pathogen like SARS-CoV-2, she may be more likely to take precautions such as avoiding indoor dining or wearing an N95 mask. Theorem 3 suggests that in the presence of such complementarity, a data provider concerned about privacy should adopt a geometric mechanism when publishing information. But in many other settings — such as those where data users wish to take one action when the population statistic is high or low, and a different action at intermediate levels — this complementarity is absent, and the geometric mechanism may not be optimal. Instead, our characterization of the solution to the data provider's problem is limited to the more general statements of Theorem 1, and its counterpart for oblivious mechanisms, Corollary 2.

To establish our third main result (Theorem 3), we rely on a partial order on information structures about ordered states of the world, and show that it ranks signals higher when they are more useful to decision makers with supermodular problems. This ordering is novel: Because of the nature of the data provider's constrained information design problem, we cannot appeal to existing tools for ranking experiments according to their usefulness to a class of decision makers.[5] Instead, we introduce the *Uniform-Peaked Relative Risk Order (UPRR)*, which ranks information structures higher when the posteriors they

[4]For 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), while 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).

[5]In particular, the information structures satisfying Corollary 2'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).

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 suboptimal 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 ϵ-geometric mechanism is UPRR-dominant over *all* ϵ-differentially private data publication mechanisms (Lemma 2), Theorem 3 follows.

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 approach also relates to prior work in computer science by Ghosh et al. (2012) and Brenner and Nissim (2014), who consider the circumstances in which specific publication mechanisms are optimal. In a similar setting to ours, Ghosh et al. (2012) show that when data users minimize a symmetric loss function of the true population statistic, the ϵ-geometric mechanism outperforms any other oblivious mechanism. When, in addition, data users have priors about the population statistic but view the database as ambiguous, and are ambiguity-averse in the maxmin sense (Gilboa et al., 1989), they show the geometric mechanism is optimal more generally.

Our Theorem 3 arrives at a related conclusion: the geometric mechanism is optimal among oblivious mechanisms in *any* setting where higher actions and higher population statistics are complementary. When respondents are anonymous, the geometric mechanism's superiority in the supermodular case extends beyond oblivious mechanisms to *all* ϵ-differentially private data publication mechanisms, even when data users are expected utility maximizers (as in our model). In contrast, Brenner and Nissim (2014) show that the optimality results in Ghosh et al. (2012) do not extend to data publication problems beyond publishing population counts for a single characteristic. We give a different counterexample—Example 3—that shows the geometric mechanism's optimality does not

extend to other important settings where data users' payoffs are not supermodular.

This paper follows other recent work that uses the tools of information design to study privacy. Ichihashi (2020) considers the way that consumers will choose to flexibly disclose information about their valuations to sellers in an online marketplace. There, privacy is an endogenous choice by the individuals described in the data, rather than an exogenous constraint to protect those agents' information from being disclosed by a third party (as in our setting). In a related setting with a single seller, Hidir and Vellodi (2021) consider the consumer-optimal design of privacy regulation — which takes the form of an information structure about the consumer's value — subject to incentive compatibility for the consumer. This approach is conceptually related to ours, since differential privacy can be viewed as an approximate incentive compatibility constraint (McSherry and Talwar, 2007); however, the data provider in our model works to benefit agents who use the data, rather than those whose data is used. In addition, our characterization of the data provider's problem is 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).

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.[6]

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 (Theorem 1). Section 4 asks when this problem's dimensionality can be reduced by restricting attention to oblivious mechanisms, and shows that we can do

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

so when respondents are anonymous (Theorem 2). Finally, Section 5 gives our optimality result for supermodular problems (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 are N *respondents*, who each have a *type* $\theta_n \in \{0,1\}$; a data provider, or *designer*, who chooses a mechanism for releasing information about their types; and a a *decision maker*, who uses that information to make better decisions.[7] The designer and decision maker each have the same prior belief $\pi_0 \in \Delta(\Theta)$ about the profile $\theta = (\theta_1, \ldots, \theta_N) \in \Theta \equiv \{0,1\}^N$ of respondents' types, or *database*, with supp $\pi_0 = \Theta$.[8,9]

Data Publication and Differential Privacy

Before observing the respondents' types, the designer chooses a *data publication mechanism* (S, m), where S is some countable set of outputs and $m : \Theta \to \Delta(S)$ maps each database to a distribution over outputs in S. 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*.

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 $n \in \{1, \ldots, N\}$, $\theta_{-n} = \theta'_{-n}$), we have[10]

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

[7] We assume a single decision maker for simplicity: All of our conclusions are unchanged if the designer seeks to maximize the sum of multiple decision makers' payoffs, so long as each agent's payoff only depends on their own action and the state.

[8] For 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 $\Delta(S)$.

[9] One 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 an individual'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.

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

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.[11] When an agent views types as independent *a priori*, this is equivalent to a limit on the amount that observing the mechanism's output can cause an agent to shift her beliefs about an individual respondent's type θ_n, regardless of how much information that agent 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 count of type-1 respondents; the permutation of those types is not payoff-relevant. Hence, his Bernoulli payoffs can be written as a function $u : A \times \Omega \to \mathbb{R}$ of his action a and the *state* $\omega_\theta \equiv \sum_{n=1}^{N} \theta_n$, where $\Omega \equiv \{0, \ldots, N\}$; 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

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

To make the problem facing the designer concrete, consider the following example, which we return to throughout the paper.

Example 1 (COVID-19 Testing)**.** A university has tested two members of its economics department for COVID-19 as part of its asymptomatic surveillance testing program. It would like to disclose the department's positivity rate to other faculty — even those who are not close contacts of those being tested — so that they can make more informed decisions about how much to interact with their colleagues. However, the university is also concerned that doing so could compromise the privacy of the faculty who have been tested, by allowing

[11] Because 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.

someone to make an inference about their status. To this end, the university has given the tested faculty (the $N = 2$ respondents) a guarantee that their test results (their types θ_n, with $\theta_n = 1$ denoting a positive result) will be ϵ-differentially private for $\epsilon = 1$.

Since the department's faculty work in the same building, the surveillance test results are correlated. In particular, we have $\pi_0((0,0)) = \pi_0((1,1)) = \frac{1}{3}$, while $\pi_0((0,1)) = \pi_0((1,0)) = \frac{1}{6}$.

The decision maker — another faculty member — must decide whether to go to lunch with a group of less COVID-averse colleagues (action x) or not (action z). None of these colleagues are among those tested as part of the surveillance program. Not going to lunch always yields this decision maker a payoff of 1. Going to lunch, on the other hand, gives the decision maker a payoff that is decreasing in ω, the number of positive tests in the department:

$$u(z,\omega) = 1 \text{ for each } \omega; \qquad u(x,\omega) = \begin{cases} 3, & \omega = 0; \\ 0, & \omega = 1; \\ -2.5, & \omega = 2. \end{cases}$$

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. The data provider is also under an obligation to guarantee differential privacy at a fixed level. Our model corresponds more precisely to any case where the data provider will publish the number of people in some category — 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.[12]

Our model does not describe how data should optimally be collected. Instead, it takes as given the data collection strategy. This allows us to focus on the relationship between privacy protection and the usefulness of published data, but means we do not account for the possible interplay between data collection and privacy protection. In some settings, a data provider could decide ahead of time to counterbalance some of the noise induced by privacy protection by collecting a larger sample. Our model could be extended to understand such decisions, but they are excluded from consideration here.

Like most of the formal privacy literature, we assume the true value of the population

[12]Note 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.

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.

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,

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

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

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

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 5), 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.[13] Figure 1 illustrates in the context of Example 1.

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

[13] It 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.

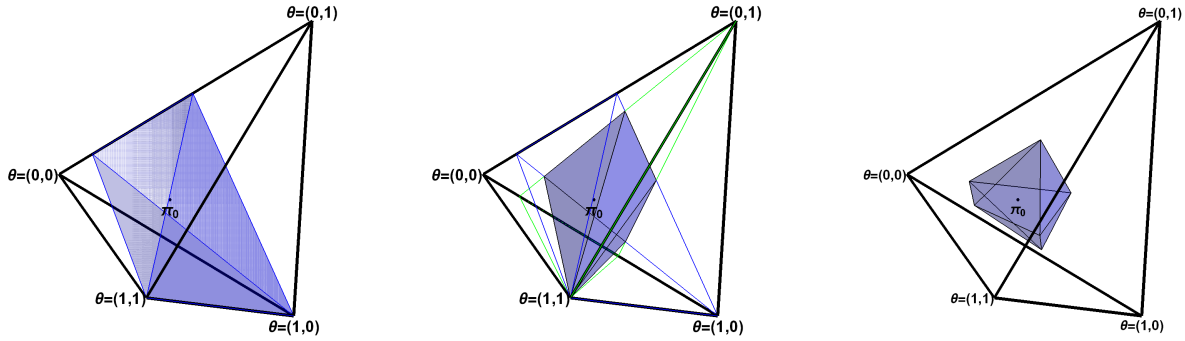


Figure 1: ϵ-differentially private posteriors in Example 1. Recall that in Example 1, $N = 2$, $\epsilon = 1$, and $(\pi_0(0,0), \pi_0(0,1), \pi_0(1,0), \pi_0(1,1)) = (\frac{1}{3}, \frac{1}{6}, \frac{1}{6}, \frac{1}{3})$. Left panel: The blue shaded region of the probability simplex is described by the constraint $-\epsilon \leq \log(\mu((0,1))/\mu((0,0))) - \log(\mu_0((0,1))/\mu_0((0,0))) \leq \epsilon$ 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 constraint $-\epsilon \leq \log(\mu((1,0))/\mu((0,0))) - \log(\mu_0((1,0))/\mu_0((1,0))) \leq \epsilon$ bounding the amount of information that an ϵ-differentially private mechanism can reveal 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.

choosing a Bayes-plausible distribution $\nu \in \Delta(\Delta(\Theta))$[14] 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*

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

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 *vertices* of the polyhedron $K(\epsilon, \pi_0)$.

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

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 $v_K : K \to \mathbb{R} \cup \{-\infty\}$. Following Kamenica and Gentzkow (2011), define the *K-restricted concavification* $V_K : K \to \mathbb{R}$ of v_K as the smallest concave function that lies above the designer's value function v on the set of posteriors $K \subseteq \Delta(\Theta)$.[15]

Theorem 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 $\pi \in \operatorname{supp} \nu^*$ *is an extreme point of* $K(\epsilon, \pi_0)$*;*

(b) Each $\pi \in \operatorname{supp} \nu^*$ *achieves the privacy bound between at least* $2^N - 1$ *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* $2^N - 1$ *distinct combinations* (θ, θ') *that have* $\theta_{-n} = \theta'_{-n}$ *for some* n*;*

(c) The support of ν^* *is a linearly independent set of vectors in* $\mathbb{R}^{2^N}$*; and*[16]

(d) ν^* *is the unique Bayes-plausible distribution of posteriors with support* $\operatorname{supp} \nu^*$.

Theorem 1 shows 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: data providers can focus on mechanisms that induce posteriors that attain the privacy bound between at least $2^N - 1$ pairs of adjacent databases (iib).

For intuition for (iia), recall that the Blackwell (1953) informativeness order ranks one signal higher than another if it is more useful to any decision maker, or equivalently if the less informative signal can be generated from the more informative one through *garbling*, or composition with another signal. If a distribution on $K_\Omega(\epsilon, \mu_0)$ induces posteriors which are not extreme points, it follows that it can be generated by garbling some distribution on the vertices of $K_\Omega(\epsilon, \mu_0)$ — and so is dominated in the Blackwell order.

For (iib), recall that the extreme points of a polyhedron in $\mathbb{R}^L$ defined by inequality constraints — such as $K(\epsilon, \pi_0)$, for $L = 2^N$ — are precisely those points where L linearly independent constraints bind. In the case of $K_\Omega(\epsilon, \mu_0)$, one of these must be the constraint that the posterior probabilities sum to 1; (iib) follows.

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

[16] For concreteness, let each belief $\pi \in \Delta(\Theta)$ be represented by the vector in $\mathbb{R}^{2^N}$ whose nth entry corresponds to the probability $\pi(\theta)$ it places on the database θ that is the binary number for n.

For (iic) and (iid), Carathéodory's theorem allows us to restrict attention to distributions of posteriors with affinely independent support (Lipnowski and Mathevet, 2017). We build on this conclusion by making two observations. First, every affinely independent set of posteriors is linearly independent (Lemma 7 in the appendix), implying (iic). This is because posteriors are elements of the simplex $\Delta(\Theta)$, which lies in an affine space of dimension 2^N. Second, the Bayes-plausible distribution supported by a linearly independent set of posteriors is unique, implying (iid). Since the set of extreme points of $K(\epsilon, \pi_0)$ — and thus the collection of its linearly independent subsets — is finite, it follows that the infinite set of ϵ-differentially private publication mechanisms can be reduced to a finite subset that are possibly optimal.[17]

4 Oblivious Mechanisms

Our results in Section 3 show 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 $2^N - 1$, the number of possible databases minus one.

But because the decision maker's payoff only depends on the state 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 state. That is, letting $P : \Delta(\Theta) \to \Delta(\Omega)$ be the projection operator defined by $P\pi(\omega) = \sum_{\theta:\omega_\theta=\omega} \pi(\theta)$, we have

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

Figure 2 illustrates.

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 state, which (following the formal privacy literature) we refer to as *oblivious mechanisms*. 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 state $\Delta(\Omega)$ directly, rather than choosing a distribution of posteriors about the database and projecting them onto $\Delta(\Omega)$.

It turns out that in general, oblivious mechanisms are not necessarily without loss. (See Example 2 for a counterexample where every oblivious mechanism is suboptimal.) But when respondents are anonymous, in the sense that the prior over databases is symmetric,

[17] More specifically, the set of publication mechanisms that are possibly optimal is finite up to inducing the same distribution of posteriors.

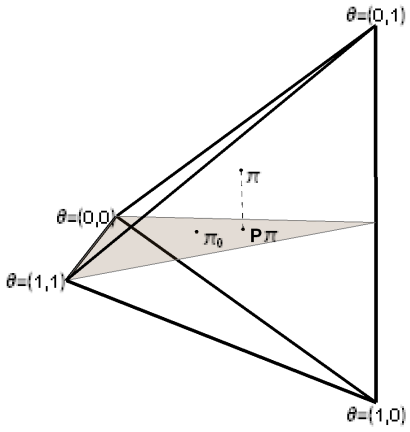


Figure 2: Projection onto $\Delta(\Omega)$ in Example 1. A posterior belief $\pi \in \Delta(\Theta)$ about the database θ is projected onto the space $\Delta(\Omega)$ — shown here embedded in $\Delta(\Theta)$ — of beliefs about the state ω.

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: In most cases, firms and statistical agencies that face the design problem considered in this paper choose to adopt mechanisms that depend only on the state. As we show, selecting a mechanism from this class is optimal whenever the designer and data users view respondents as interchangeable, but not necessarily otherwise. 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.

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 $\omega_\theta = \omega_{\theta'}$.[18,19] Hence, there is a function $\sigma : \Omega \to \Delta(S)$ 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 *Laplace*, *Gaussian*, and *geometric* mechanisms are each oblivious, since they publish the sum of the true state ω_θ and a random variable.

[18]Note that in our setting, this is equivalent to the requirement that $m(\cdot|\theta) = m(\cdot|\theta')$ whenever θ is a permutation of θ'.

[19]In the terminology of the computer science literature on formal privacy, the state ω_θ is the outcome of a *counting query* applied to θ, and our definition specifies that a mechanism is oblivious with respect to that query.

When a mechanism is oblivious, Proposition 2 shows that the differential privacy criterion (1) simplifies to a limit on the amount that that *moving to an adjacent state* (i.e., from $\omega - 1$ to ω or vice versa) can change the distribution of the mechanism's realizations. Intuitively, if a pair of databases are equal in all but one entry, then their counts of respondents with type $\theta_n = 1$ — their states — must differ by exactly one. Consequently, differential privacy's restriction on posterior beliefs about the *database* (4) reduces to a restriction on the posterior beliefs that the mechanism induces about the *state*.

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. $\left|\log\left(\frac{\sigma(s|\omega)}{\sigma(s|\omega-1)}\right)\right| \leq \epsilon$ *for each $s \in S$ and $\omega \in \{1, \ldots, N\}$.*

iii. For each posterior belief about the state $\mu \in \Delta(\Omega)$ induced by (S, σ),

$$\left|\log\left(\frac{\mu(\omega)}{\mu(\omega-1)}\right) - \log\left(\frac{\mu_0(\omega)}{\mu_0(\omega-1)}\right)\right| \leq \epsilon \text{ for each } \omega \in \{1, \ldots, N\}. \tag{7}$$

We call the set of posterior beliefs about the state that satisfy (7) the *oblivious ϵ-differentially private posteriors* $K_\Omega(\epsilon, \mu_0) \subset \Delta(\Omega)$. Figure 3 illustrates.

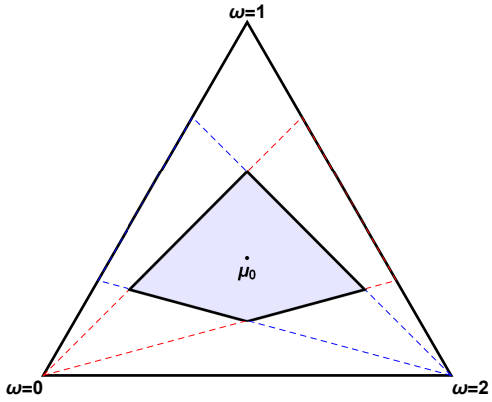


Figure 3: Oblivious ϵ-differentially private posteriors in Example 1. Recall that in Example 1, $N = 2$, $\epsilon = 1$, and $\mu_0 = \left[\frac{1}{3}\ \frac{1}{3}\ \frac{1}{3}\right]'$. The region of the probability simplex enclosed by the blue dotted lines is described by the constraint $-\epsilon \leq \log(\mu(1)/\mu(0)) - \log(\mu_0(1)/\mu_0(0)) \leq \epsilon$; 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 state* by a non-oblivious mechanism. In general, the designer can induce a distribution τ of posterior beliefs about the state 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.[20] The set of distributions that satisfy this criterion coincides with those that can be induced with an 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 8 in the appendix).

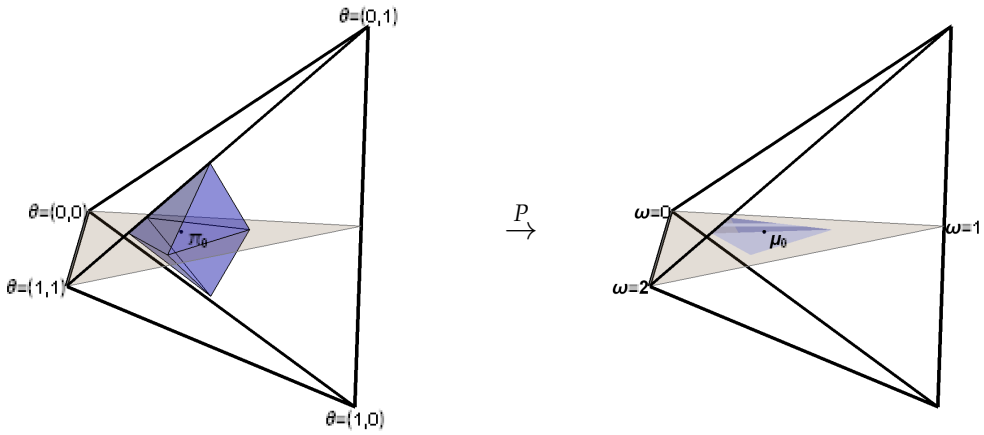


Figure 4: Projection of $K(\epsilon, \pi_0)$ onto $\Delta(\Omega)$ in Example 1. Left panel: The set of ϵ-differentially private posteriors (blue) and the space $\Delta(\Omega)$ of posteriors about the state (brown). Right panel: Since π_0 is symmetric in Example 1, the projection operator P carries $K(\epsilon, \pi_0)$ to the set of oblivious ϵ-differentially private posteriors (blue).

The main result of this section shows that this always holds — and thus oblivious mechanisms are without loss — when *respondents are anonymous*, in the sense that the prior distribution π_0 is symmetric: $\pi_0(\theta) = \pi_0(\theta')$ whenever θ is a permutation of θ'.[21]

Theorem 2 (Anonymity and Oblivious Mechanisms)**.** *If respondents are anonymous, then for each ϵ-differentially private data publication mechanism (S, m), there exists an ϵ-differentially pri-*

[20] That 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.

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

vate oblivious data publication mechanism (S, σ) *that induces the same distribution of posterior beliefs about the state.*

As a consequence of Theorem 2, whenever respondents are anonymous, 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 1. *If respondents are anonymous, the oblivious mechanism* (S, σ) *solves the designer's problem (2) if and only if it induces a distribution of posteriors about the state that solves*

$$\max_{\tau \in \Delta(K_\Omega(\epsilon, \mu_0))} \{E_\tau \hat{v}(\mu) \text{ s.t. } E_\tau \mu = \mu_0\}. \tag{8}$$

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, databases with the same state must be permutations of each other. Hence, a symmetric prior places equal probability on each of them. But equally importantly, each differs by a single entry from the same number of databases whose state is lower by 1.[22] As a consequence of these two 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 $N = 2$, the set of databases that differ from θ in one entry is *invariant* under permutation.

Proposition 3 (Oblivious Mechanisms with Two Respondents)**.** *If* $N = 2$*, then for each* ϵ*-differentially private data publication mechanism* (S, m)*, there exists an* ϵ*-differentially private oblivious data publication mechanism* (S, σ) *that induces the same distribution of posterior beliefs about the state.*

The conclusion of Theorem 2 is markedly different from that of Ghosh et al. (2012), who also consider a model with binary data. When the decision maker seeks to minimize a function of the distance between their action and the state,[23] they show that oblivious

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

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

mechanisms are without loss if the decision maker has a prior belief μ_0 about the state 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 on whether certain respondents are *a priori* more likely than others to have the characteristic in question. As Example 2 illustrates, absent such symmetry, there may be a non-oblivious mechanism that outperforms every oblivious ϵ-differentially private mechanism.

Example 2 (Generality of Oblivious Mechanisms: A Counterexample)**.** Consider the COVID-19 testing setting from Example 1, but now suppose that another person in the economics department has been tested, so that $N = 3$. Furthermore, this third faculty member is respondent 2's spouse, and so their test results are highly correlated. Specifically, conditional on the first two faculty members' results (θ_1, θ_2), the probability that $\theta_3 = \theta_2$ is $1 - \delta$ for some small $\delta > 0$; i.e., we have

$$\pi_0((0,0,0)) = \pi_0((1,1,1)) = (1-\delta)/3 \quad \text{and} \quad \pi_0((0,1,1)) = \pi_0((1,0,0)) = (1-\delta)/6,$$
$$\text{but } \pi_0((0,0,1)) = \pi_0((1,1,0)) = \delta/3 \quad \text{and} \quad \pi_0((0,1,0)) = \pi_0((1,0,1)) = \delta/6.$$

For small enough δ, a non-oblivious mechanism can induce posteriors about the state that are inaccessible to ϵ-differentially private oblivious mechanisms without changing the level of privacy loss. In particular, in the limit $\delta \to 0$, a state of $\omega = 1$ *always* corresponds to the database $(1,0,0)$, while a state of $\omega = 2$ *always* corresponds to the database $(0,1,1)$. But these databases differ in *all three* entries, so differential privacy only indirectly restricts the amount that the posterior probability of one can differ from the posterior probability of the other. This allows the designer to provide much more information about whether the state ω is 1 rather than 2.

Specifically, writing each $\pi \in \Delta(\{0,1\}^3)$ as the vector

$$\begin{bmatrix}\pi((0,0,0)) & \pi((1,0,0)) & \pi((0,1,0)) & \pi((0,0,1)) & \pi((1,1,0)) & \pi((1,0,1)) & \pi((0,1,1)) & \pi((1,1,1))\end{bmatrix}',$$

consider the posterior

$$\hat{\pi} = \frac{\phi \circ \pi_0}{\phi \cdot \pi_0}, \text{ where } \phi = \begin{bmatrix}e^{-2\epsilon} & e^{-3\epsilon} & e^{-\epsilon} & e^{-\epsilon} & e^{-2\epsilon} & e^{-2\epsilon} & 1 & e^{-\epsilon}\end{bmatrix}',$$

where $\circ$ denotes the elementwise (Hadamard) product. This posterior is ϵ-differentially private: $\hat{\pi} \in K(\epsilon, \pi_0)$. It achieves the upper privacy bound $\frac{\pi((1,0,0))/\pi_0((1,0,0))}{\pi((0,0,0))/\pi_0((0,0,0))} = e^{\epsilon}$ between $(0,0,0)$ and $(1,0,0)$, and the lower privacy bound $\frac{\pi((1,1,1))/\pi_0((1,1,1))}{\pi((0,1,1))/\pi_0((0,1,1))} = e^{-\epsilon}$ between $(0,1,1)$ and $(1,1,1)$, while achieving the privacy bounds between other databases in a way that dis-

tinguishes between $(1,0,0)$ and $(0,1,1)$ as much as possible.[24] As $\delta \to 0$, its third through sixth entries vanish along with the corresponding entries of π_0. Hence, its projection $P\hat{\pi}$ onto $\Delta(\Omega)$ approaches

$$\hat{\mu} = \frac{\psi \circ \mu_0}{\psi \cdot \mu_0}, \text{ where } \psi = \begin{bmatrix} e^{-\epsilon} & e^{-3\epsilon} & 1 & e^{-2\epsilon} \end{bmatrix}'.$$

This posterior about the state is outside of $K_\Omega(\epsilon, \mu_0)$, and so cannot be induced with an oblivious mechanism: it exceeds the upper privacy bound $\frac{\mu(2)/\mu_0(2)}{\mu(1)/\mu_0(1)} \leq e^\epsilon$ between states $\omega = 1$ and $\omega = 2$. Consequently, when δ is small enough, and θ_2 and θ_3 are very highly correlated, oblivious mechanisms are not always optimal.[25]

4.3 Information Design with Oblivious Mechanisms

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

Corollary 2 (Characterization of Optimal Data Publication Mechanisms)**.** *Suppose that respondents are anonymous.*

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 state $\tau^* \in \Delta(\Delta(\Omega))$ *such that*

(a) Each $\mu \in \operatorname{supp} \tau^*$ *is an extreme point of* $K_\Omega(\epsilon, \mu_0)$*;*

(b) Each $\mu \in \operatorname{supp} \tau^*$ *achieves the privacy bound at each state:* $\left|\log\left(\frac{\mu(\omega)}{\mu(\omega-1)}\right) - \log\left(\frac{\mu_0(\omega)}{\mu_0(\omega-1)}\right)\right| = \epsilon$ *for all* $\mu \in \operatorname{supp} \tau^*$ *and* $\omega \in \Omega \setminus \{0\}$*;*

(c) The support of τ^* *is a linearly independent set of vectors in* $\mathbb{R}^{N+1}$*; and*

[24] Specifically, it achieves the upper privacy bound between $(1,0,0)$ and both $(1,1,0)$ and $(1,0,1)$, between $(0,0,0)$ and both $(0,1,0)$ and $(0,0,1)$, between $(0,1,0)$ and $(0,1,1)$, between $(0,1,0)$ and $(0,1,1)$, between $(1,1,0)$ and $(1,1,1)$, and between $(1,0,1)$ and $(1,1,1)$; and the lower privacy bound between $(0,1,0)$ and $(1,1,0)$ and between $(0,0,1)$ and $(1,0,1)$.

[25] Consider, for instance, a decision maker who takes action z when his belief about the state is in $K_\Omega(\epsilon, \mu_0)$, but takes a different action, x, when his belief about the state is in some neighborhood of $\hat{\mu}$. (To see how this might occur, suppose that distinguishing between $\omega = 1$ and $\omega = 2$ is important for the decision maker's choice, but distinguishing between the other states is not, e.g., because x gives a much worse payoff than z when $\omega = 1$, a somewhat better payoff when $\omega = 2$, and the same payoff when $\omega \in \{0, 3\}$.) Then any oblivious ϵ-differentially private mechanism does not offer any useful information to the decision maker, but a non-oblivious ϵ-differentially private mechanism that induces $\hat{\pi}$ does.

(d) τ^* is the unique Bayes-plausible distribution of posteriors with support supp τ^*.

Example 1 revisited. In the simple COVID-19 surveillance testing setting considered in Example 1, oblivious mechanisms are without loss, both since the prior π_0 is symmetric (Theorem 2) and since $N = 2$ (Proposition 3). Then we have $\hat{v}_{K_\Omega(\epsilon,\mu_0)}(\mu) = \max\left\{1, \begin{bmatrix} 3 \\ 0 \\ 2.5 \end{bmatrix} \cdot \mu\right\}$, as shown in Figure 5, below. By Corollary 2, the maximized value of the designer's problem is given by its concavification $\hat{V}_{K_\Omega(\epsilon,\mu_0)}$ (left panel, in blue) evaluated at μ_0 (right panel). Moreover, this value can be achieved with a distribution supported by a linearly independent set of extreme points of $K_\Omega(\epsilon, \mu_0)$. With $N + 1 = 3$, there are only two of these: the unique Bayes-plausible distribution supported by the extreme points of $K_\Omega(\epsilon, \mu_0)$ that achieve the upper privacy bound $\frac{\mu(\omega)/\mu(\omega-1)}{\mu_0(\omega)/\mu_0(\omega-1)} = e^\epsilon$ for both state $\omega = 1$ and state $\omega = 2$, neither state 1 nor state 2, and for state 1 only, respectively; and the unique Bayes-plausible distribution supported by the extreme points of $K_\Omega(\epsilon, \mu_0)$ that achieve the upper privacy bound for state 1 only and state 2 only, respectively. As the right panel of Figure 5 shows, the former is optimal in Example 1.

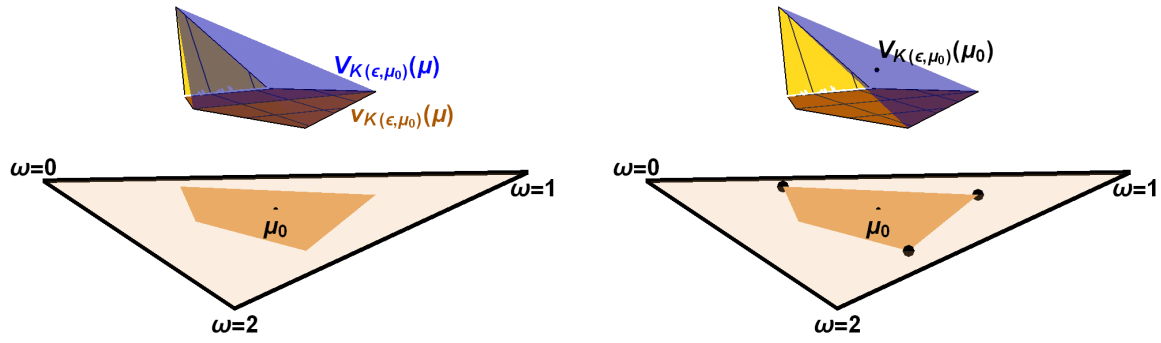


Figure 5: Optimal data publication mechanisms in Example 1. Left panel: $\hat{v}_{K_\Omega(\epsilon,\mu_0)}$ (orange graph) and its concavification $\hat{V}_{K_\Omega(\epsilon,\mu_0)}$ (blue) on the set of ϵ-differentially private posteriors (shaded region of simplex). Right panel: The problem's maximized value $\hat{V}_{K_\Omega(\epsilon,\mu_0)}(\mu_0)$ and the posteriors that support its solution, labeled by their signatures.

The characteristics (iib)-(iid) of the optimal oblivious mechanism from Corollary 2 allow us to pin down its functional form and connect it to an important mechanism from the differential privacy literature. Given a finite set S, a *query function* $q : \{0,1\}^N \times S \to \mathbb{R}$, and a *base measure* $\xi \in \Delta(S)$, define the *exponential mechanism* (McSherry and Talwar, 2007) as the oblivious mechanism $(S, \mathcal{E}^\epsilon_{q,\xi})$, where

$$\mathcal{E}^\epsilon_{q,\xi}(s|\theta) \equiv \frac{\exp(\epsilon q(\theta, s))\xi(s)}{\sum_{r\in S}\exp(\epsilon q(\theta, r))\xi(r)}. \tag{9}$$

When we allow the base measure to depend on the privacy loss — i.e., when we take a *tuple* of base measures $\{\xi_x\}_{x\in(0,\infty)}$ as given, and replace ξ with ξ_ϵ in (9) — we call $(S, \mathcal{E}^\epsilon_{q,\xi_\epsilon})$

a *generalized exponential mechanism*. As Corollary 3 shows, the oblivious data publication mechanisms described in Corollary 2 (ii) are always members of a class of generalized exponential mechanisms parameterized by the sets of privacy bounds (upper or lower) attained by the posterior beliefs about the state that the mechanism induces. Consequently, if respondents are anonymous, then for any $\epsilon > 0$, one of these generalized exponential mechanisms solves the designer's problem (2).

Corollary 3 (Optimality of Generalized Exponential Mechanisms)**.** *There is a solution to the designer's problem (2) that belongs to a family of oblivious generalized exponential mechanisms* $\{(\{1,\dots,|\mathcal{P}|\}, \mathcal{E}^{\epsilon}_{q_{\mathcal{P}},\xi_{\epsilon,\mathcal{P}}})\}_{\mathcal{P}\subseteq 2^{\Omega\setminus\{0\}}}$ *parameterized by the collection* $\mathcal{P} \subseteq 2^{\Omega\setminus\{0\}}$ *of sets of states* $\Phi_{\mu_j} \subseteq 2^{\Omega\setminus\{0\}}$ *where the posteriors* μ_j *that the mechanism induces attain the upper privacy bound.*

Since Corollary 2's optimal mechanism induces a distribution of posteriors about the state that has linearly independent support, that distribution can be expressed as the projection of the prior about the state μ_0 on the posteriors in the distribution's support. And since each of those posteriors μ attains the privacy bound at each state (Corollary 2 (iib)), their forms are pinned down by the set Φ_μ of states where they achieve the *upper* privacy bound $\frac{\mu(\omega)/\mu(\omega-1)}{\mu_0(\omega)/\mu_0(\omega-1)} = e^{\epsilon}$ (Lemma 10 in the appendix). Together, these facts pin down the mechanism's form, and parameterize it by the collection $\mathcal{P}$ of sets of states where the posteriors it induces attain the upper privacy bound.

5 Supermodular Payoffs and the Geometric Mechanism

In a large class of applications where respondents are anonymous, we can move beyond Corollary 2's characterization, and show that a specific oblivious publication mechanism solves the designer's problem. Observe that the optimal mechanism in Example 1 induced beliefs that satisfied the upper privacy bound at each state below some (posterior-specific) cutoff x, and the lower privacy bound at each state above the cutoff. That is, it induced precisely those posteriors that achieve the upper privacy bound on *lower sets* of states of the form $(\Omega \setminus \{0\}) \cap (-\infty, x]$. As we show in Lemma 1, this property characterizes the distribution of posteriors about the state induced by the well-known *ϵ-geometric mechanism* $(\mathbb{Z}, \sigma^{g}_{\epsilon})$ (Ghosh et al., 2012), which adds two-sided geometrically-distributed noise to the true state, and so has the signal distribution[26]

$$\sigma^{g}_{\epsilon}(s|\omega) = \left(\frac{1-e^{-\epsilon}}{1+e^{-\epsilon}}\right) e^{-\epsilon|s-\omega|}.$$

[26]While the ϵ-geometric mechanism produces more outputs than there are states — which, by Corollary 2 (iic), is unnecessary — it induces the same distribution of posterior beliefs about the state as the *truncated*

In this section, we show that if respondents are anonymous, the ϵ-geometric mechanism solves the designer's problem whenever each of the decision makers have totally ordered actions and supermodular payoffs (Theorem 3). Because of 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 order on information structures — the Uniform-Peaked Relative Risk Order — 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-higher information structure is superior in any supermodular decision problem with actions on the real line (Theorem 4).

5.1 Optimality of the Geometric Mechanism

In a key contribution to the privacy literature, Ghosh et al. (2012) show that when $A_i = \Omega$ for each i, and decision makers aim to "match the state", i.e.,

$$u_i(a_i, \omega) = \tilde{u}_i(|a_i - \omega|, \omega) \text{ for } \tilde{u}_i \text{ nonincreasing in } |a_i - \omega|,$$

the ϵ-geometric mechanism is optimal among all ϵ-differentially private oblivious mechanisms. This result has led to the adoption of the geometric mechanism (or its continuous analogue, the Laplace mechanism) by statistical agencies such as the U.S. Census Bureau (Foote et al., 2021). However, its scope — and thus the class of environments where the geometric mechanism is appropriate — is unclear.[27] On the one hand, the environment of Example 1 is not captured by the "match the state" setting of Ghosh et al. (2012), suggesting that the geometric mechanism is optimal in a broader class of environments. On the other, in many settings outside the class considered by Ghosh et al. (2012), the geometric mechanism is suboptimal. Example 3 illustrates.

Example 3 (Optimality of the Geometric Mechanism: A Counterexample)**.** Suppose that the university from Example 1 has once again performed COVID-19 tests on its economics

ϵ-geometric mechanism $(\Omega, \hat{\sigma}^g_\epsilon)$ (Ghosh et al., 2012) with signal distribution

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

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.

[27] For instance, Brenner and Nissim (2014) show that the results of Ghosh et al. (2012) do not generalize to settings where (in our terminology) respondents' types are arbitrary integers, where they are multidimensional, or when decision makers' payoffs depend on the database in ways that are not captured by the count of type-1 respondents.

faculty. However, instead of conducting tests for active infection, the university has instead tested for levels of neutralizing antibodies above some threshold (whether acquired from vaccination or prior infection), and the two faculty members tested comprise our decision maker's lunch group.

Once again, the prior over test results is symmetric, and correlated so as to induce a uniform prior about the state: $\pi_0((0,0)) = \pi_0((1,1)) = \frac{1}{3}$, while $\pi_0((0,1)) = \pi_0((1,0)) = \frac{1}{6}$, and so $\mu_0 = [\frac{1}{3}\ \frac{1}{3}\ \frac{1}{3}]'$. The difference between this setting and that of Example 1 instead lies in the decision maker's payoff from going to lunch: Now the decision maker's payoff is higher if either $\omega = 2$ (in which case the rest of her lunch group has enough immunity that infection is unlikely) or $\omega = 0$ (in which case she may infer that COVID-19 has not recently spread among her colleagues) than if $\omega = 1$ (in which case she may infer that COVID-19 may be actively spreading among faculty). Specifically,

$$u(z,\omega) = 1; \qquad u(x,\omega) = \begin{cases} 2.5, & \omega \in \{0,2\}; \\ -2.5, & \omega = 1. \end{cases}$$

Consequently, $\hat{v}_{K_\Omega(\epsilon,\mu_0)}(\mu) = \max\left\{1, \begin{bmatrix} 2.5 \\ -2.5 \\ 2.5 \end{bmatrix} \cdot \mu\right\}$.

The left panel of Figure 6 illustrates the value of the ϵ-geometric mechanism as the height at μ_0 of the plane through the points on the graph of $v_{K_\Omega(\epsilon,\mu_0)}$ corresponding to the posteriors about the state that the mechanism induces. Here, as Lemma 1 shows, those are the extreme points of $K_\Omega(\epsilon,\mu_0)$ that achieve the upper privacy bound for neither state 1 nor state 2, for state 1 only, and for both state 1 and state 2. The right panel of Figure 6, on the other hand, illustrates the strictly larger value of the oblivious mechanism that is optimal in this example.

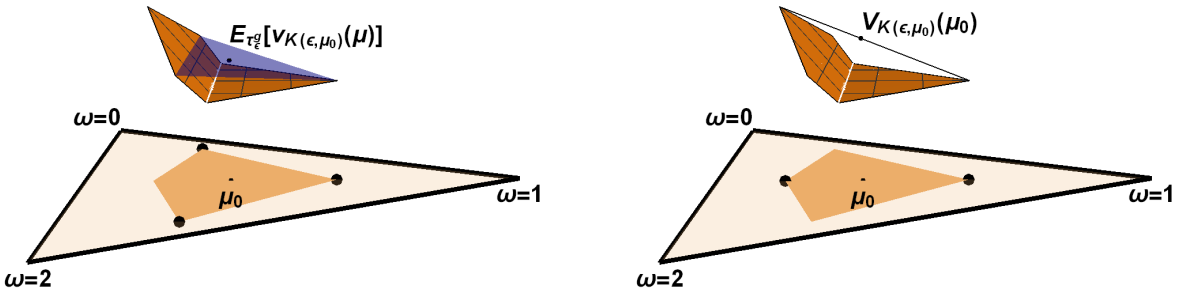


Figure 6: The geometric and optimal data publication mechanisms in Example 3. Both panels: The set of oblivious ϵ-differentially private posteriors $K_\Omega(\epsilon,\mu_0)$ (shaded region of simplex), and the restriction of the designer's value function to it $\hat{v}_{K_\Omega(\epsilon,\mu_0)}$ (orange graph). Left panel: Theposterior beliefs about the state induced by the ϵ-geometric mechanism, the hyperplane (blue) through the points on the graph of $\hat{v}_{K_\Omega(\epsilon,\mu_0)}$ corresponding to those posteriors, and the value of the ϵ-geometric mechanism, $E_{\tau_\epsilon^g}[\hat{v}_{K_\Omega(\epsilon,\mu_0)}(\mu)]$. Right panel: The maximized value $\hat{V}_{K_\Omega(\epsilon,\mu_0)}(\mu_0)$ of the designer's problem (8) and the posterior beliefs about the state that support its solution.

To help understand why the ϵ-geometric mechanism is optimal in Example 1 and the setting of Ghosh et al. (2012), but not in Example 3, we begin by characterizing the distribution τ_ϵ^g of posterior beliefs about the state that it induces. Lemma 1 shows that τ_ϵ^g is supported by precisely those extremal ϵ-differentially private posteriors that are *unimodally* different from the prior — i.e., those that achieve the upper privacy bound on a lower set.

Lemma 1 (Posteriors Produced by the Geometric Mechanism)**.** *The ϵ-geometric mechanism induces precisely those 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 $x > 0$:*

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

To understand Lemma 1, observe that if a decision maker's belief about the state 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.[28]

Since they achieve the upper privacy bound for states $\omega \leq x$ and the lower privacy bound for states $\omega > x$, the posteriors induced by each output x of the geometric mechanism are collapsed around state x (or the state closest to x, if $x \notin [0, N]$) as much as the differential privacy constraint will allow. This is intuitively optimal when decision makers want to minimize a nonincreasing function of the distance between their action and the state of the world, as in Ghosh et al. (2012), since it allows decision makers to be more certain that their chosen action is near the true state. But equally intuitively, it is suboptimal when decision makers are more interested in whether the state is central or extreme — as in Example 3.

Theorem 3 shows that the crucial feature that leads to the geometric mechanism's optimality is not that the decision makers want to *match* the state, but rather that they want to take *higher actions in higher states*. In particular, whenever each decision maker's actions can be ordered from highest to lowest (e.g., from least to most precaution, as in Example 1), and each decision maker has supermodular payoffs — that is, their marginal benefit of taking a higher action is increasing in the state[29] — the ϵ-geometric mechanism outperforms

[28] Note 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.

[29] Recall that when $A_i \subseteq \mathbb{R}$, $u_i : A_i \times \Omega \to \mathbb{R}$ is supermodular if for each $a \leq a'$ and each $\omega \leq \omega'$,

any other ϵ-differentially private oblivious mechanism.

Theorem 3 (Optimality of the Geometric Mechanism for Supermodular Problems)**.** *Suppose that respondents are anonymous. 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 The UPRR Order on 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 the optimal mechanism from Example 3 — are not even ordered by first-order stochastic dominance.[30] This makes results like those of Lehmann (1988), Quah and Strulovici (2009), and Athey and Levin (2018) 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 order on information structures in which the ϵ-geometric data publication mechanism dominates all other ϵ-differentially private oblivious data publication mechanisms.

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

In words, one distribution of posteriors 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.[31] Equivalently, the *relative risk* $\mu(\omega)/\mu'(\omega)$ of a state ω under a posterior μ in the dominant distri-

$u_i(a', \omega') - u_i(a, \omega') \geq u_i(a', \omega) - u_i(a, \omega)$.

[30] Recall 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 $(S, \sigma : \Omega \to \Delta(S))$, 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).

[31] That is, if the ratio of the probability that the former posterior places on a state closer to the its peak to the probability it places on a state further away from its peak is at least the ratio of probabilities placed on those states by the latter posterior.

bution'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.[32]

Example 4. The UPRR order allows us to compare the distribution of posteriors about the state τ^* that was optimal in Example 3 with the distribution τ_ϵ^g induced by the ϵ-geometric mechanism. The former is supported by the two extreme points of $K_\Omega(\epsilon, \mu_0)$ that achieve exactly one upper privacy bound (at either state 1 or state 2) The latter, on the other hand, induces the extreme points of $K_\Omega(\epsilon, \mu_0)$ that achieve the upper privacy bound for both states, neither state, and for state 1 only, respectively.

Below, in Figure 7, we graph all four posterior beliefs $\mu \in \text{ext}(K_\Omega(\epsilon, \mu_0))$ that achieve the privacy bound at each state, as well as the relative risk $\mu(\omega)/\mu'(\omega)$ for each $\mu \in \text{supp}\, \tau_\epsilon^g$ and each $\mu' \in \text{supp}\, \tau^*$. One can see that for each $\mu \in \text{supp}\, \tau_\epsilon^g$, the relative risk against either posterior $\mu' \in \text{supp}\, \tau^*$ has the same peak $\omega^*(\mu)$: when μ attains both lower privacy bounds, $\omega^*(\mu) = 0$; when μ attains the upper privacy bound for state 1 and the lower privacy bound for state 2, $\omega^*(\mu) = 1$; and when μ attains both upper privacy bounds, $\omega^*(\mu) = 2$. Hence, $\tau_\epsilon^g \succeq_{UPRR} \tau^*$.

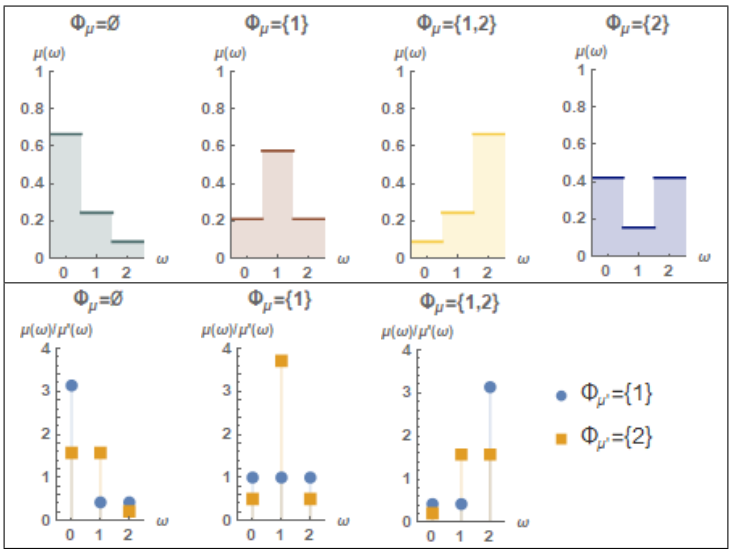


Figure 7: Ordering the distributions induced by the ϵ-geometric mechanism from Examples 1 and 3 and the optimal data publication mechanism from Example 3 in the UPRR order. Top panel: extreme points of $K_\Omega(\epsilon, \mu_0)$, identified by the set of states Φ_μ at which they attain the upper privacy bound. Bottom panel: Relative risk of each state under beliefs $\mu \in \text{supp}\, \tau_\epsilon^g$ versus beliefs $\mu' \in \text{supp}\, \tau^*$.

[32]Note 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 $[\hat{\omega}^*(s), \infty)$.

Lemma 2 shows that the UPRR-dominance of the ϵ-geometric mechanism demonstrated in Example 4 extends to *all* ϵ-differentially private data publication mechanisms with any number of respondents.

Lemma 2 (Geometric Mechanisms are UPRR-Dominant)**.** *If $\tau \in \Delta(\Delta(\Omega))$ is induced by an ϵ-differentially private oblivious mechanism, then $\tau_\epsilon^g \succeq_{UPRR} \tau$, where for each $\mu \in \operatorname{supp} \tau_\epsilon^g$, $\omega^*(\mu)$ 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 state at or below some x and the lower privacy bound at each higher state. Hence, these posteriors each concentrate probability around their peak $\omega^*(\mu) = x$ as much as the differential privacy constraint will allow. This is precisely what is necessary for the distribution they support to UPRR-dominate the distributions of posteriors induced by every other differentially private mechanism.

We conclude with our general comparative statics result on the UPRR order. Theorem 4 shows that among information structures which induce finitely many posteriors, those that are higher in this order 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 2), and the designer's problem can be solved by inducing a distribution of posteriors with finite support (Corollary 2 (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 $\tau \succeq_{UPRR} \tau'$, then for any compact $A \subseteq \mathbb{R}$ and continuous supermodular function $h : \Omega \times A \to \mathbb{R}$, $E_\tau[\max_{a \in A} E_\mu[h(\omega, a)]] \geq E_{\tau'}[\max_{a \in A} E_\mu[h(\omega, a)]]$.*

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 4).[33]

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, Theorem 1 and Corollary 2 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 only depend on the true value of the population statistic that data users are interested in. We show that using a mechanism from this class is optimal in the very common case where respondents are anonymous (Theorem 2) — for instance, if the data are i.i.d. — but not necessarily otherwise (Example 2). Finally, we show that among oblivious mechanisms, the commonly used *geometric mechanism* is optimal whenever data users view their actions and the number of respondents with the characteristic of interest as complementary, in the sense that their payoffs are supermodular.

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

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

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 Theorem 1 can be given in such settings, it is less clear whether oblivious mechanisms are without loss under a symmetry condition, as in Theorem 2. We leave the question of when oblivious mechanisms are without loss in such settings to future work.

Finally, in many settings, a designer wishes to publish information about the population average of a characteristic that is quantitative (i.e., takes integer values) rather than qualitative (i.e., takes binary values, as in this paper). There, Brenner and Nissim (2014) have shown that in the "match the state" setting of Ghosh et al. (2012), the geometric mechanism is no longer optimal among oblivious mechanisms. However, since our comparative static (Theorem 4) applies to any setting where states are totally ordered — including when they represent population averages of a quantitative characteristic — it may facilitate an optimality result for such settings that is similar to Theorem 3, but with a different mechanism being optimal.

Oblivious mechanisms may not be without loss in settings where respondent characteristics take integer values, though, even with a symmetry assumption. A key observation used in our characterization of posteriors about the state that can be induced with an oblivious ϵ-differentially private mechanism (Proposition 2) is that changing one respondent's data always changes the state by the same amount. This is not true for averages of quantitative data: the amount that changing one respondent's data can change the population average is dependent not just on the value of the population average, but on the way that the characteristic is distributed in the sample. Hence, if *some* database with population average ω is adjacent to *some* database with population average ω', then obliviousness imposes the privacy bound between those two databases on *every* pair of databases with population averages ω and ω', even those where the database with population average ω is not adjacent to *any* database with population average ω'. Consequently, working with oblivious mechanisms meaningfully increases the restrictiveness of differential privacy.[34]

However, just as in settings with multiple characteristics, it is straightforward to show that when respondent characteristics take integer values, the designer's problem still has

[34]Specifically, unlike with binary data, adding together the bounds that differential privacy places on a non-oblivious mechanism's likelihood ratios of adjacent databases no longer allows us to recover the bounds that it places on the likelihood ratios of population averages under an oblivious mechanism, even when the prior is symmetric: $PK(\epsilon, \pi_0) \neq K_\Omega(\epsilon, \mu_0)$.

a characterization analogous to Theorem 1. This suggests that the information design approach we take in this paper may offer new insights in such environments — insights which must differ in important ways from those we give here.

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.

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.

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.

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.

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.

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.

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.

LIPNOWSKI, E. AND L. MATHEVET (2017): "Simplifying Bayesian Persuasion," Working paper.

MATYSKOVÁ, L. (2019): "Bayesian Persuasion With Costly Information Acquisition," Working paper.

MCSHERRY, F. AND K. TALWAR (2007): "Mechanism Design via Differential Privacy," in *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, IEEE, 94–103.

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.

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. KANCHA, 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.

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. (2021): "Designing Incentives for Heterogeneous Researchers," Working paper.

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 $U([0,1])$.[35,36] Next, in Lemma 3, we describe the sense in which this representation is equivalent to the original distribution. Finally, Proposition 4 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,

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

since (W, X) is equal in distribution to $(\omega, R_\tau(j))$ on the latter vector's support.[37] 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 $\mathcal{J} : \operatorname{supp}\tau \to \mathbb{R}$ 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

[35] In 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.

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

[37] We 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.

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 5 illustrates the Frechét representation concept in the context of Examples 1 and 3.

Example 5 (Frechét Representations of the Geometric Mechanism)**.** Consider the ϵ-geometric mechanism $(\mathbb{Z}, \sigma_\epsilon^g)$ in the simple setting from Example 1. 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 8, and does so according to the distribution σ_ϵ^g displayed in the upper-left of Figure 8. By Lemma 2, 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 8.

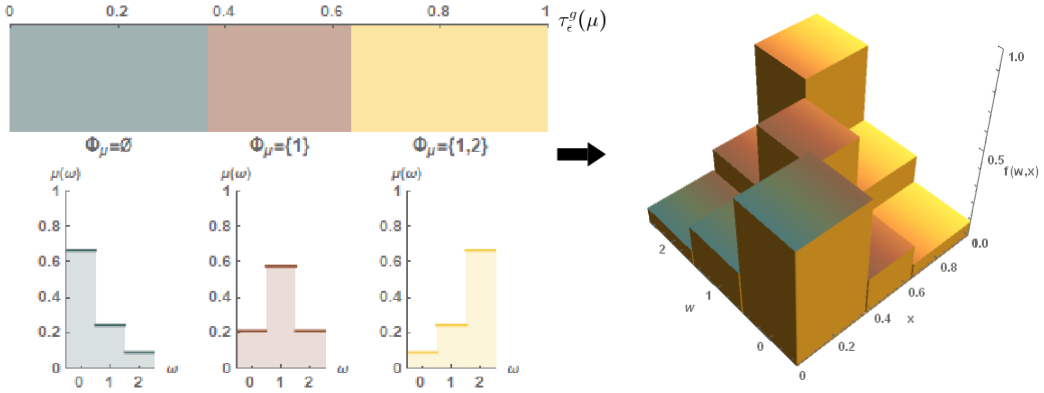


Figure 8: Generating the Frechét representation of τ_ϵ^g with respect to the peak ω^* in Examples 1 and 3. Upper-left panel: The distribution τ_ϵ^g. Lower-left panel: The posterior beliefs in supp τ_ϵ^g. Right panel: The mixed density of the Frechét representation of τ_ϵ^g with respect to ω^*.

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 $t : \text{supp}\,\tau \to \mathbb{R}$, 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 $F(w,x) = \int_0^x \sum_{y=0}^w f(y,z)dz$, 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 5, continued. Consider once more the geometric mechanism from Example 1. This time, let $t : \operatorname{supp} \tau_\epsilon^g \to \mathbb{R}$ 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 9) "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.

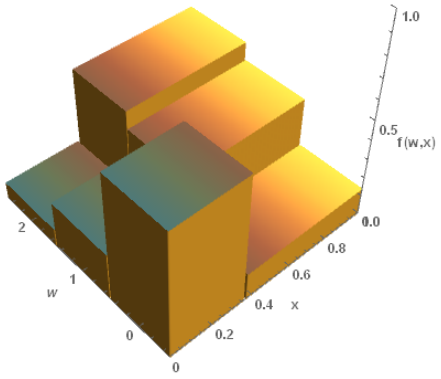


Figure 9: The Frechét representation of τ_ϵ^g with respect to t in Examples 1 and 3. For each $\omega \in \Omega$, representing τ_ϵ^g with respect to t instead of ω^* averages the representation's mixed density $f(w,x)$ over the region of $\{\omega\} \times [0,1]$ associated with posteriors that attain 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 (The-

orem 4), we employ a Frechét representation with respect to an endogenous function.[38]

Lemma 3 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 3 (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 $t : \operatorname{supp} \tau \to A$, and let F be the cdf of (W, X).*

i. $F \in \mathcal{M}(\mu_0, U([0,1]))$.

ii. For any measurable function $h : \Omega \times A \to \mathbb{R}$,

$$E_\tau[E_\mu[h(\omega, t(\mu))]] = E[h(W, Q_{\tau,t}(X))].$$

iii. For any upper semicontinuous function $h : \Omega \times A \to \mathbb{R}$, any $\tau' \in \Delta(\Delta(\Omega))$ with finite support, and any $s : \operatorname{supp} \tau' \to A$,

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

Parts (ii) and (iii) of Lemma 3 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 Xth percentile value of $t(\mu)$ when μ is distributed according to τ — h's second argument is instead $Q_{\tau',s}(X)$, the Xth 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

[38] Specifically, we use a Frechét representation with respect to a selection from the decision maker's solution correspondence.

respect to X would.[39] 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 $\mu \in \operatorname{supp}\tau$. (In fact, it may do strictly worse, since — as in Figure 9 — 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 ith elements $\mathbf{X}_i$ and $\mathbf{Y}_i$ take values in a compact set $S_i \subseteq \mathbb{R}$, and for any supermodular function $h : \prod S_i \to \mathbb{R}$, 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 4 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 3 then yields Theorem 4.

Proposition 4 (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 $\omega^* : \operatorname{supp}\tau \to \mathbb{R}$. 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 $t : \operatorname{supp}\tau' \to \mathbb{R}$ in the supermodular stochastic order: for any supermodular function $h : \Omega \times [0,1] \to \mathbb{R}$,*

$$E[h(W,X)] \geq E[h(\hat{W},\hat{X})].$$

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 4 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

[39] That 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, $E[h(W,a^*(X))] \geq E[h(W,Q_{\tau',s}(X))]$.

$\bar{F}(w,x) \geq \bar{G}(w,x)$ 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 4 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 $\mu \in \operatorname{supp} \tau$ 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 $\mu' \in \operatorname{supp} \tau'$; 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 4 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 Proofs[40]

Lemma 4. *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* $\operatorname{supp} \nu \subset K(\epsilon,\pi_0)$.

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 $\operatorname{supp} \nu \subset K(\epsilon,\pi_0)$ as well. □

Proof of Proposition 1 (Differentially Private Data Publication as Information Design) Follows immediately from Lemma 4. □

Lemma 5 (Characterization of Differentially Private Posteriors)**.** *For each $\epsilon > 0$,*

[40] For a set S, denote its relative interior by $\operatorname{ri}(S)$ and its relative boundary by $\operatorname{rbd}(S)$. 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 ith entry if $i \in S$ and 0 in its ith 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 and $\mathbb{1}$ to denote an indicator function. 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 $\mathbb{R}^{N+1}$ whose nth entry corresponds to the probability $\mu(n)$ it places on state n; i.e., the indices of the vector's entries and the states they represent are the same. Likewise, we can let each belief $\pi \in \Delta(\Theta)$ be represented by the vector in $\mathbb{R}^{2^N}$ whose nth entry corresponds to the probability $\pi(\theta)$ it places on the database θ that is the binary number for n. (For simplicity, we write $\mathbf{1}_\theta \in \mathbb{R}^{2^N}$ for the indicator vector for the index whose binary number is θ.)

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* $\mathbb{R}^{2^N}$.

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 $\theta''_{-n} = \theta'_{-n}$ for some n, 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 $\theta_{-n} = \theta'_{-n}$ for some n, 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,

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

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

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

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

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

(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. □

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 5. □

Proof of Proposition 2 (Differential Privacy for Oblivious Mechanisms) ((i)⇒(ii)) For each $\omega \in \Omega \setminus \{0\}$, choose $n \in \{1, \dots, N\}$, and choose θ such that $\theta_n = 1$ and $\omega_\theta = \omega$, and choose θ' such that $\theta_{-n} = \theta'_{-n}$ and $\theta_n = 0$. Then since (S, σ) is ϵ-differentially private, for each $s \in S$, $|\log(\sigma(s|\omega)/\sigma(s|\omega - 1))| = |\log(\sigma(s|\omega_\theta)/\sigma(s|\omega_{\theta'}))| \leq \epsilon$, since $\omega_{\theta'} = \omega - 1$; (ii) follows.

((ii)⇒(i)) If $\theta, \theta' \in \{0,1\}^N$ are such that $\theta_{-n} = \theta'_{-n}$ for some n, then either $\theta = \theta'$, in which case (1) holds trivially, or $|\omega_\theta - \omega_{\theta'}| = 1$, 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)⇔(iii)) Follows from Bayes' rule, since

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

Lemma 6 (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* $\operatorname{supp}\hat{\nu} \subseteq \operatorname{ext}(K(\epsilon, \pi_0))$ *that is more informative than ν (in the Blackwell sense).*

ii. *Each extreme point of $K(\epsilon, \pi_0)$ attains at least $2^N - 1$ privacy bounds: For each $\pi \in \operatorname{ext}(K(\epsilon, \pi_0))$, $\left|\log\left(\frac{\pi(\theta)}{\pi(\theta')}\right) - \log\left(\frac{\pi_0(\theta)}{\pi_0(\theta')}\right)\right| = \epsilon$ for at least $2^N - 1$ distinct combinations (θ, θ') that have $\theta_{-n} = \theta'_{-n}$ for some n.*

Proof. (i): By Lemma 5 (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 $K(\epsilon, \pi_0) = \operatorname{conv}(\operatorname{ext}(K(\epsilon, \pi_0)))$. Then for each $\pi \in \operatorname{supp}\nu$, we can write $\pi = \sum_{\pi' \in \operatorname{ext}(K(\epsilon,\pi_0))} \lambda(\pi, \pi')\pi'$ for some $\{\lambda(\pi, \pi')\}_{\pi' \in \operatorname{ext}(K_\Omega(\epsilon,\mu_0))} \subset [0,1]$ with $\sum_{\pi' \in \operatorname{ext}(K(\epsilon,\pi_0))} \lambda(\pi, \pi') = 1$. 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 $P = \{x | x \cdot z_i \geq b_i, i \in \{1, \dots, d\}\}$ for $\{z_i\}_{i=1}^d \subset \mathbb{R}^L$ and $\{b_i\}_{i=1}^d \subset \mathbb{R}$ are precisely the basic feasible solutions of a linear program on P; i.e., those $x \in P$ such that $x \cdot z_i = b_i$ for a linearly independent subset of the z_i with L elements. By 5 (iii), $K(\epsilon, \pi_0) \subset \mathbb{R}^{2^N}$ is a closed convex polyhedron described by the inequalities in (10). It follows that at least 2^N 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 5 (i), and so for any $\theta, \theta' \in \Theta$, we cannot have both

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

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 $2^N - 1$ distinct combinations (θ, θ') that have $\theta_{-n} = \theta'_{-n}$ for some n. □

Corollary 5 (Characterization of Extreme Points of $K_\Omega(\epsilon, \mu_0)$)**.**

i. If the support of a distribution ν 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 state $\hat{\tau} \in \Delta(\Delta(\Omega))$ with $\operatorname{supp} \hat{\tau} \subseteq \operatorname{ext}(K_\Omega(\epsilon, \mu_0))$ 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: $\operatorname{ext}(K_\Omega(\epsilon, \mu_0)) = \left\{\mu \in \Delta(\Omega) \mid \left|\log\left(\frac{\mu(\omega)}{\mu(\omega-1)}\right) - \log\left(\frac{\mu_0(\omega)}{\mu_0(\omega-1)}\right)\right| = \epsilon \ \forall \omega \in \Omega \setminus \{0\}\right\}$.

Proof. (i): Follows identically to Lemma 6 (i).

(ii): Recall that the extreme points of a polyhedron $P \subset \mathbb{R}^{N+1}$ defined by $P = \{\mu | \mu \cdot z_i \geq b_i, i \in \{1, \ldots, d\}\}$ 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 5 (iii), $K_\Omega(\epsilon, \mu_0)$ is a closed convex polyhedron described by the inequalities

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

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 5 (i), and so for any $\omega \in \Omega \setminus \{0\}$, we cannot have both

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

Hence, if $\mu \in \operatorname{ext}(K_\Omega(\epsilon, \mu_0))$, 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 $Z = \{\mathbf{1}\} \cup \{e^{\epsilon}\frac{\mu_0(\omega)}{\mu_0(\omega-1)}\mathbf{1}_{\omega-1}\omega - \mathbf{1}_\omega | \omega \in S\} \cup \{e^{-\epsilon}\frac{\mu_0(\omega)}{\mu_0(\omega-1)}\mathbf{1}_{\omega-1} - \mathbf{1}_\omega | \omega \in (\Omega \setminus \{0\}) \setminus S\}$ 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 $\mu \in \operatorname{ext}(K_\Omega(\epsilon, \mu_0))$. □

Lemma 7 (Distributions of Posteriors with Affine and Linear Independent Supports)**.**

i. If $\nu \in \Delta(\Delta(\Theta))$ is Bayes-plausible and $\operatorname{supp} \nu$ is finite,[41] *there is a Bayes-plausible $\nu' \in$*

[41] The finiteness assumption is unnecessary here, but is adopted for simplicity.

$\Delta(\Delta(\Theta))$ *such that the posteriors in* $\operatorname{supp}\nu'$ *are affinely independent and* $\operatorname{supp}\nu' \subseteq \operatorname{supp}\nu$.

ii. If $\{\pi_j\}_{j=1}^J \subset \Delta(\Theta)$ *is affinely independent, it is linearly independent.*

iii. If $\nu \in \Delta(\Delta(\Theta))$ *is Bayes-plausible and* $\operatorname{supp}\nu$ *is linearly independent, then* $\operatorname{supp}\nu' \neq \operatorname{supp}\nu$ *for any Bayes-plausible* $\nu' \in \Delta(\Delta(\Theta)), \nu' \neq \nu$.

Proof. (i): Since $\operatorname{supp}\nu$ is finite and ν is Bayes-plausible, we have $\sum_{\pi\in\operatorname{supp}\nu}\nu(\pi)\pi = \pi_0$, and so $\pi_0 \in \operatorname{conv}(\operatorname{supp}\nu)$. Then by Carathéodory's theorem (e.g., Rockafellar (1997) Theorem 17.1) there is some affinely independent $\{\pi_j\}_{j=1}^J \subseteq \operatorname{supp}\nu$ such that $\pi_0 \in \operatorname{conv}(\{\pi_j\}_{j=1}^J)$. 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 $0 = \sum_{\theta\in\Theta}\sum_{j=1}^J \lambda_j\pi_j(\theta) = \sum_{j=1}^J \lambda_j \sum_{\theta\in\Theta}\pi_j(\omega)$. 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 $\operatorname{supp}\nu' = \operatorname{supp}\nu$ for some Bayes-plausible $\nu' \neq \nu$. Since $\operatorname{supp}\nu$ is linearly independent, it must be finite. Since ν and ν' are both Bayes-plausible, we must have $\sum_{\pi\in\operatorname{supp}\nu}\nu(\pi)\pi = \sum_{\pi\in\operatorname{supp}\nu}\tau'(\pi)\pi = \pi_0$. Then we have $\sum_{\pi\in\operatorname{supp}\nu}(\nu(\mu) - \nu'(\mu))\pi = 0$, with $\nu(\pi) - \nu'(\pi) \neq 0$ for some $\pi \in \operatorname{supp}\nu$, since $\nu \neq \nu'$. Then $\operatorname{supp}\nu$ is linearly dependent, a contradiction. □

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

i. If $\tau \in \Delta(\Delta(\Omega))$ *is Bayes-plausible and* $\operatorname{supp}\tau$ *is finite, there is a Bayes-plausible* $\tau' \in \Delta(\Delta(\Omega))$ *such that the posteriors in* $\operatorname{supp}\tau'$ *are affinely independent and* $\operatorname{supp}\tau' \subseteq \operatorname{supp}\tau$.

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* $\operatorname{supp}\tau$ *is linearly independent, then* $\operatorname{supp}\tau' \neq \operatorname{supp}\tau$ *for any Bayes-plausible* $\tau' \in \Delta(\Delta(\Omega)), \tau' \neq \tau$.

Proof. Follows identically to Lemma 7. □

Proof of Theorem 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 5 (ii), $\pi_0 \in \operatorname{ri}(K(\epsilon,\pi_0))$. Since $K(\epsilon,\pi_0) \subseteq \Delta(\Theta)$, $K(\epsilon,\pi_0)$ is bounded in $\mathbb{R}^{2^N}$; then by Lemma 5 (iii), it is compact and convex. (i) then follows from Kamenica and Gentzkow (2011) Online Appendix Proposition 3.

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 $\operatorname{supp}\nu \not\subseteq \operatorname{ext}(K(\epsilon,\pi_0))$, then by Lemma 6 (i) we can choose a $\hat{\nu}$ with $\operatorname{supp}\hat{\nu} \subseteq \operatorname{ext}(K(\epsilon,\pi_0))$ which is Blackwell-more informative than ν and hence also solves (8). If $\operatorname{supp}\nu \subseteq \operatorname{ext}(K(\epsilon,\pi_0))$, choose $\hat{\nu} = \nu$. Then $\operatorname{supp}\hat{\nu}$ 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 7 (i), there is a Bayes-plausible ν^* such that $\operatorname{supp}\nu^*$ is an affinely independent set and $\operatorname{supp}\nu^* \subseteq \operatorname{supp}\hat{\nu} \subseteq \operatorname{ext}(K(\epsilon,\pi_0))$. Moreover, by Yoder (2021) Lemma 3, $V_{K(\epsilon,\pi_0)}(\pi) = v_{K(\epsilon,\pi_0)}(\pi) = v(\pi)$ for each $\pi \in \operatorname{supp}\hat{\nu}$, and in addition, $V_{K(\epsilon,\pi_0)}$ is affine on $\operatorname{conv}(\operatorname{supp}\hat{\nu})$: there exists $x \in \mathbb{R}^{2^N}$ such that for each $\pi \in \operatorname{conv}(\operatorname{supp}\hat{\nu}))$, $V_{K(\epsilon,\pi_0)}(\pi) = V_{K(\epsilon,\pi_0)}(\pi_0) + x \cdot (\pi - \pi_0)$. Then since $\operatorname{supp}\nu^* \subseteq \operatorname{supp}\hat{\nu}$, 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 $(\operatorname{supp}\nu^*, m^*)$ 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 $(\operatorname{supp}\nu^*, m^*)$ 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, $(\operatorname{supp}\nu^*, m^*)$ solves the designer's problem (2). Since it induces ν^*, (iic) follows from affine independence of $\operatorname{supp}\nu^*$ and Lemma 7 (ii); (iid) then follows from Lemma 7 (iii). And since $\operatorname{supp}\nu^* \subseteq \operatorname{ext}(K(\epsilon,\pi_0))$, (iib) follows from Lemma 6 (i). □

Lemma 8. *The following are equivalent:*

i. If the distribution $\tau \in \Delta(\Delta(\Omega))$ of posterior beliefs about the state 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)⇒(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 4, $K_\Omega(\epsilon,\mu_0) \subseteq PK(\epsilon,\pi_0)$. 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 5 (ii), there exists $\delta > 0$ such that $\pi' = \pi_0 - \delta(\pi - \pi_0) \in K(\epsilon,\pi_0)$. Then the distribution ν with

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

is Bayes-plausible with $\operatorname{supp}\nu \subset K(\epsilon,\pi_0)$, and so by Lemma 4 can be induced by an ϵ-differentially private mechanism, which thus induces the distribution of posteriors about the state $\tau = \nu \circ P^{-1} \in \Delta(\Delta(\Omega))$. Then $P\pi \in \operatorname{supp}\tau$, 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)⇒(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 state induced by the ϵ-differentially private mechanism (S,m). Then $\operatorname{supp}\tau = P\operatorname{supp}\nu$, where ν is the distribution of posterior beliefs about θ induced by (S,m). Since (S,m) is ϵ-differentially private, by Lemma 4, $\operatorname{supp}\nu \subset K(\epsilon,\pi_0)$. Then $\operatorname{supp}\tau = P\operatorname{supp}\nu \subset PK(\epsilon,\pi_0) = K_\Omega(\epsilon,\mu_0)$, and by Proposition 2, τ can be induced by an oblivious ϵ-differentially private mechanism. □

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

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

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

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

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

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

$$e^{-\epsilon} \sum_{\theta':\omega_{\theta'}=\omega} \pi(\theta') \frac{\mu_0(\omega-1)}{\mu_0(\omega)} \leq P\pi(\omega-1) \leq e^{\epsilon} \sum_{\theta':\omega_{\theta'}=\omega} \pi(\theta') \frac{\mu_0(\omega-1)}{\mu_0(\omega)}$$

$$e^{-\epsilon} P\pi(\omega) \frac{\mu_0(\omega-1)}{\mu_0(\omega)} \leq P\pi(\omega-1) \leq e^{\epsilon} P\pi(\omega) \frac{\mu_0(\omega-1)}{\mu_0(\omega)},$$

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 4, $K_\Omega(\epsilon, \mu_0) \subseteq PK(\epsilon, \pi_0)$. So $PK(\epsilon, \pi_0) = K_\Omega(\epsilon, \mu_0)$; the statement follows by Lemma 8. □

Proof of Corollary 1 (Differentially Private Data Publication as Information Design) Follows immediately from (6), Proposition 2, and Theorem 2. □

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

$$\frac{P\pi(2)}{\mu_0(2)} = \frac{\pi((1,1))}{\pi_0((1,1))}$$

$$e^{-\epsilon}\pi((0,1)) \leq \pi((1,1)) \frac{\pi_0((0,1))}{\pi_0((1,1))} \leq e^{\epsilon}\pi((0,1))$$

$$e^{-\epsilon}\pi((1,0)) \leq \pi((1,1)) \frac{\pi_0((1,0))}{\pi_0((1,1))} \leq e^{\epsilon}\pi((1,0))$$

$$\Rightarrow e^{-\epsilon}(\pi((1,0)) + \pi((0,1))) \leq \pi((1,1)) \frac{\pi_0((1,0)) + \pi_0((0,1))}{\pi_0((1,1))} \leq e^{\epsilon}(\pi((1,0)) + \pi((0,1)))$$

$$\Rightarrow e^{-\epsilon} P\pi(1) \leq P\pi(2) \frac{\pi_0(1)}{\pi_0(2)} \leq e^{\epsilon} P\pi(1).$$

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 4, $K_\Omega(\epsilon, \mu_0) \subseteq PK(\epsilon, \pi_0)$. Thus, $PK(\epsilon, \pi_0) = K_\Omega(\epsilon, \mu_0)$; the statement follows by Lemma 8. □

Lemma 9. *If a distribution $\tau^* \in \Delta(\Delta(\Omega))$ is Bayes-plausible, and its support is linearly independent, it can be written as the projection of the prior μ_0 on τ^*'s support: Labeling* $\operatorname{supp} \tau^* = \{\mu_j\}_{j=1}^J$, *and letting* $M_{\tau^*} \equiv \begin{bmatrix} \mu_1 & \cdots & \mu_J \end{bmatrix}$ *be the matrix whose columns are the posteriors in*

$\operatorname{supp} \tau^*$, *we have*[42]

$$\begin{bmatrix} \tau^*(\mu_1) & \cdots & \tau^*(\mu_J) \end{bmatrix}' = (M'_{\tau^*} M_{\tau^*})^{-1} M'_{\tau^*} \mu_0.$$

Proof. Let $\vec{\tau}^* = \begin{bmatrix} \tau^*(\mu_1) & \dots & \tau^*(\mu_J) \end{bmatrix}'$, and note that since τ^* is Bayes-plausible, $M_{\tau^*} \vec{\tau}^* = \mu_0$. By Corollary 2 (iic), the columns of M_{τ^*} are linearly independent, and so M_{τ^*} has a left inverse given by $(M'_{\tau^*} M_{\tau^*})^{-1} M'_{\tau^*}$, and we have $\vec{\tau}^* = (M'_{\tau^*} M_{\tau^*})^{-1} M'_{\tau^*} M_{\tau^*} \vec{\tau}^* = (M'_{\tau^*} M_{\tau^*})^{-1} M'_{\tau^*} \mu_0$, as desired. □

Lemma 10 (Extremal Posteriors and Upper Privacy Bounds). *For each* $\mu \in \operatorname{ext}(K_\Omega(\epsilon, \mu_0))$,

$$\mu = \frac{\mu_0 \circ \psi(\Phi_\mu, \epsilon)}{\mu_0 \cdot \psi(\Phi_\mu, \epsilon)},$$

where $\circ$ *denotes the entrywise (Hadamard) product,* $\Phi_\mu \equiv \left\{ \omega \in \Omega \setminus \{0\} \mid \frac{\mu(\omega)/\mu(\omega-1)}{\mu_0(\omega)/\mu_0(\omega-1)} = e^\epsilon \right\}$ *is the set of states at which* μ *attains the upper privacy bound, and* $\psi(\Phi, \epsilon) \in \mathbb{R}^{N+1}$ *is defined entrywise by*[43]

$$\psi(\Phi, \epsilon)_\omega \equiv \exp\left(\epsilon \sum_{i=1}^{\omega} (\mathbb{1}\{i \in \Phi\} - \mathbb{1}\{i \notin \Phi\}) \right).$$

Proof. Let $\mu \in \operatorname{ext}(K_\Omega(\epsilon, \mu_0))$. From the definition of Φ_μ, for each $\omega \in [0, N]$, we have

$$\mu(\omega) = \frac{\mu_0(\omega)}{\mu_0(\omega-1)} \mu(\omega-1) \exp\left(\epsilon (\mathbb{1}\{\omega \in \Phi\} - \mathbb{1}\{\omega \notin \Phi\}) \right).$$

Recursively substituting for $\mu(\omega - i)$, $i \in \{1, \dots, \omega\}$, yields

$$\mu(\omega) = \frac{\mu_0(\omega)}{\mu_0(0)} \mu(0) \exp\left(\epsilon \sum_{i=1}^{\omega} (\mathbb{1}\{i \in \Phi\} - \mathbb{1}\{i \notin \Phi\}) \right), \tag{14}$$

and so for each $\omega \neq 0$, $\mu(\omega)$ is just $\frac{\mu_0(\omega)}{\mu_0(0)} \mu(0)$ times the ωth entry of $\psi(\Phi_\mu, \epsilon)$. Since $(0, 0] = \emptyset$, this also holds for $\omega = 0$. Then we have

$$\mu = \frac{\mu(0)}{\mu_0(0)} \mu_0 \circ \psi(\Phi_\mu, \epsilon). \tag{15}$$

Since $\mu \in \Delta(\Omega)$, we must have $\mu \cdot \mathbf{1} = 1$; then from (15), $1 = \mu \cdot \mathbf{1} = \frac{\mu(0)}{\mu_0(0)} \mu_0 \cdot \psi(\Phi_\mu, \epsilon)$. It

[42] Recall that we represent beliefs about the state $\mu \in \Delta(\Omega)$ as zero-indexed vectors: that is, we number the entries of μ starting from 0 rather than starting from 1, so that the entries range from 0 to N rather than 1 to $N+1$. Consequently, the entries of the matrix M_τ are zero-indexed as well.

[43] Recall that the zero-indexing convention means that the entries of $\psi(\Phi, \epsilon)$ have indices from 0 to N.

follows that $\mu(0) = \frac{\mu_0(0)}{\mu_0 \cdot \psi(\Phi_\mu, \epsilon)}$, and hence

$$\mu = \frac{\mu_0(0)}{\mu_0 \cdot \psi(\Phi_\mu, \epsilon)} \frac{1}{\mu_0(0)} \mu_0 \circ \psi(\Phi_\mu, \epsilon) = \frac{\mu_0 \circ \psi(\Phi_\mu, \epsilon)}{\mu_0 \cdot \psi(\Phi_\mu, \epsilon)},$$

as desired. □

Lemma 11 (Closed Forms for Data Publication Mechanisms). *Suppose $\tau \in \Delta(\Delta(\Omega))$ is Bayes-plausible and with linearly independent $\operatorname{supp}\tau \subseteq \operatorname{ext}(K_\Omega(\epsilon, \mu_0))$, let $J = |\operatorname{supp}\tau|$, let $\psi(\Phi, \epsilon) \in \mathbb{R}^{N+1}$ be defined entrywise by $\psi(\Phi, \epsilon)_\omega \equiv \exp\left(\epsilon \sum_{i=1}^{\omega} (\mathbb{1}\{i \in \Phi\} - \mathbb{1}\{i \notin \Phi\})\right)$, label $\operatorname{supp}\tau = \{\mu_j\}_{j=1}^J$, and let $\Psi_{\mathcal{P}_{\tau,\epsilon}} \equiv \begin{bmatrix} \psi(\Phi_{\mu_1}, \epsilon) & \cdots & \psi(\Phi_{\mu_J}, \epsilon) \end{bmatrix}$, where for each μ, $\Phi_\mu \equiv \left\{\omega \in \Omega \setminus \{0\} \mid \frac{\mu(\omega)/\mu(\omega-1)}{\mu_0(\omega)/\mu_0(\omega-1)} = e^\epsilon\right\}$ is the set of states at which μ attains the upper privacy bound.*[44] *Then $\operatorname{rank}\Psi_{\mathcal{P}_{\tau,\epsilon}} = J$, and τ is produced by the oblivious data publication mechanism $(\{1, \ldots, J\}, \sigma_{\tau,\epsilon})$, where*

$$\sigma_{\tau,\epsilon}(j|\omega) = \Sigma(\tau, \epsilon)_{\omega j}; \qquad \Sigma(\tau, \epsilon) \equiv \Psi_{\mathcal{P}_{\tau,\epsilon}} \operatorname{diag}((\Psi'_{\mathcal{P}_{\tau,\epsilon}} \Psi_{\mathcal{P}_{\tau,\epsilon}})^{-1} \Psi'_{\mathcal{P}_{\tau,\epsilon}} \mathbf{1}).$$

Proof. Let $\vec{\tau} \equiv [\tau(\mu_1) \cdots \tau(\mu_J)]'$. By Lemma 10, $M_\tau \equiv \begin{bmatrix} \mu_1 & \cdots & \mu_J \end{bmatrix} = \operatorname{diag}(\mu_0) \Psi_{\mathcal{P}_{\tau,\epsilon}} \operatorname{diag}(\Psi'_{\mathcal{P}_{\tau,\epsilon}} \mu_0)^{-1}$. Then since the columns of M_τ are linearly independent,

$$\begin{aligned} J = \operatorname{rank}(M_\tau) &\leq \min\{\operatorname{rank}(\operatorname{diag}(\Psi'_{\mathcal{P}_{\tau,\epsilon}} \mu_0)^{-1}), \operatorname{rank}(\Psi_{\mathcal{P}_{\tau,\epsilon}}), \operatorname{rank}(\operatorname{diag}(\mu_0))\} \\ &\leq \operatorname{rank}(\Psi_{\mathcal{P}_{\tau,\epsilon}}) \leq \min\{J, N+1\} = J. \end{aligned}$$

Hence, $J = \operatorname{rank}(\Psi_{\mathcal{P}_{\tau,\epsilon}})$; i.e., $\Psi_{\mathcal{P}_{\tau,\epsilon}}$ has full column rank. Then it has a left inverse given by $(\Psi'_{\mathcal{P}_{\tau,\epsilon}} \Psi_{\mathcal{P}_{\tau,\epsilon}})^{-1} \Psi'_{\mathcal{P}_{\tau,\epsilon}}$; since τ is Bayes-plausible, we have

$$\begin{aligned} \operatorname{diag}(\mu_0) \Psi_{\mathcal{P}_{\tau,\epsilon}} \operatorname{diag}(\Psi'_{\mathcal{P}_{\tau,\epsilon}} \mu_0)^{-1} \vec{\tau} = M_\tau \vec{\tau} &= \mu_0 \\ \Psi_{\mathcal{P}_{\tau,\epsilon}} \operatorname{diag}(\Psi'_{\mathcal{P}_{\tau,\epsilon}} \mu_0)^{-1} \vec{\tau} &= \operatorname{diag}(\mu_0)^{-1} \mu_0 = \mathbf{1} \\ \operatorname{diag}(\Psi'_{\mathcal{P}_{\tau,\epsilon}} \mu_0)^{-1} \vec{\tau} &= (\Psi'_{\mathcal{P}_{\tau,\epsilon}} \Psi_{\mathcal{P}_{\tau,\epsilon}})^{-1} \Psi'_{\mathcal{P}_{\tau,\epsilon}} \mathbf{1} \\ \vec{\tau} &= \operatorname{diag}(\Psi'_{\mathcal{P}_{\tau,\epsilon}} \mu_0)(\Psi'_{\mathcal{P}_{\tau,\epsilon}} \Psi_{\mathcal{P}_{\tau,\epsilon}})^{-1} \Psi'_{\mathcal{P}_{\tau,\epsilon}} \mathbf{1} \\ &= (\Psi'_{\mathcal{P}_{\tau,\epsilon}} \mu_0) \circ ((\Psi'_{\mathcal{P}_{\tau,\epsilon}} \Psi_{\mathcal{P}_{\tau,\epsilon}})^{-1} \Psi'_{\mathcal{P}_{\tau,\epsilon}} \mathbf{1}). \end{aligned} \tag{16}$$

[44] Recall that the zero-indexing convention means that the entries of each vector in Lemma 11 — and the rows of each matrix — have indices ranging from 0 to N, so that the ωth entry or row corresponds to state ω.

By Lemma 10, for each $j \in \{1, \dots, J\}$ and each $\omega \in \Omega$, we have

$$\begin{aligned}
\Sigma(\tau, \epsilon) &= \Psi_{\mathcal{P}_{\tau,\epsilon}} \operatorname{diag}(\Psi'_{\mathcal{P}_{\tau,\epsilon}} \mu_0)^{-1} \operatorname{diag}(\Psi'_{\mathcal{P}_{\tau,\epsilon}} \mu_0) \operatorname{diag}((\Psi'_{\mathcal{P}_{\tau,\epsilon}} \Psi_{\mathcal{P}_{\tau,\epsilon}})^{-1} \Psi'_{\mathcal{P}_{\tau,\epsilon}} \mathbf{1}) \\
&= \Psi_{\mathcal{P}_{\tau,\epsilon}} \operatorname{diag}(\Psi'_{\mathcal{P}_{\tau,\epsilon}} \mu_0)^{-1} \operatorname{diag}[(\Psi'_{\mathcal{P}_{\tau,\epsilon}} \mu_0) \circ ((\Psi'_{\mathcal{P}_{\tau,\epsilon}} \Psi_{\mathcal{P}_{\tau,\epsilon}})^{-1} \Psi'_{\mathcal{P}_{\tau,\epsilon}} \mathbf{1})] \\
&= \Psi_{\mathcal{P}_{\tau,\epsilon}} \operatorname{diag}(\Psi'_{\mathcal{P}_{\tau,\epsilon}} \mu_0)^{-1} \operatorname{diag}(\vec{\tau}) \qquad \text{(from (16))} \\
&= \operatorname{diag}(\mu_0)^{-1} \operatorname{diag}(\mu_0) \Psi_{\mathcal{P}_{\tau,\epsilon}} \operatorname{diag}(\Psi'_{\mathcal{P}_{\tau,\epsilon}} \mu_0)^{-1} \operatorname{diag}(\vec{\tau}) \\
&= \operatorname{diag}(\mu_0)^{-1} M_\tau \operatorname{diag}(\vec{\tau}).
\end{aligned}$$

Hence, for each $j \in \{1, \dots, J\}$ and $\omega \in \Omega$, $\sigma_{\tau,\epsilon}(j|\omega) = \frac{1}{\mu_0(\omega)} \mu_j(\omega) \vec{\tau}_j$. Then by Bayes' rule, the posterior probability placed on ω by a decision maker who observes realization j of $(\{1, \dots, J\}, \sigma_\tau)$ is given by

$$\frac{\sigma_{\tau,\epsilon}(j|\omega)\mu_0(\omega)}{\sum_{w \in \Omega} \sigma_{\tau,\epsilon}(j|w)\mu_0(w)} = \frac{\vec{\tau}_j \mu_j(\omega)}{\sum_{w \in \Omega} \vec{\tau}_j \mu_j(w)} = \frac{\mu_j(\omega)}{\sum_{w \in \Omega} \mu_j(w)} = \mu_j(\omega).$$

Hence, for each possible realization $j \in \{1, \dots, J\}$ of $(\{1, \dots, J\}, \sigma_{\tau,\epsilon})$, observing j induces the posterior μ_j. It follows that the distribution of posteriors about the state induced by $(\{1, \dots, J\}, \sigma_{\tau,\epsilon})$ has support $\{\mu_j\}_{j=1}^J = \operatorname{supp} \tau$; by Corollary 6 (iii), this distribution must be τ. □

Proof of Corollary 3 (Optimality of Exponential Mechanisms) Label $\mathcal{P} = \{\Phi_j\}_{j=1}^{|\mathcal{P}|}$, and let

$$q_{\mathcal{P}}(\theta, j) = \sum_{i=1}^{\omega_\theta} (\mathbb{1}\{i \in \Phi_j\} - \mathbb{1}\{i \notin \Phi_j\}); \qquad \xi_{x,\mathcal{P}}(j) = \frac{((\Psi'_{\mathcal{P},x} \Psi_{\mathcal{P},x})^{-1} \Psi'_{\mathcal{P},x} \mathbf{1})_j}{\mathbf{1}'(\Psi'_{\mathcal{P},x} \Psi_{\mathcal{P},x})^{-1} \Psi'_{\mathcal{P},x} \mathbf{1}},$$

where $\Psi_{\mathcal{P},\epsilon} \equiv \begin{bmatrix} \psi(\Phi_1, \epsilon) & \cdots & \psi(\Phi_{|\mathcal{P}|}, \epsilon) \end{bmatrix}$ for $\psi(\Phi, \epsilon) \in \mathbb{R}^{N+1}$ defined entrywise by $\psi(\Phi, \epsilon)_\omega \equiv \exp\left(\epsilon \sum_{i=1}^{\omega} (\mathbb{1}\{i \in \Phi\} - \mathbb{1}\{i \notin \Phi\})\right)$.

By Corollary 2 (ii), there exists an oblivious mechanism that solves the designer's problem (2) and induces a distribution of posteriors τ^* about the state such that $\operatorname{supp} \tau^*$ is linearly independent and $\operatorname{supp} \tau^* \subseteq \operatorname{ext}(K_\Omega(\epsilon, \mu_0))$; label $\operatorname{supp} \tau^* = \{\mu_j\}_{j=1}^J$. Let $\Phi_j = \Phi_{\mu_j}$ for each j. For each j, let $\Phi_j = \left\{\omega \in \Omega \setminus \{0\} \mid \frac{\mu_j(\omega)/\mu_j(\omega-1)}{\mu_0(\omega)/\mu_0(\omega-1)} = e^\epsilon\right\}$ be the set of states at which μ_j attains the upper privacy bound. Then for any $\theta \in \{0,1\}^N$, the (ω_θ, j)th element of $\Psi_{\mathcal{P},\epsilon}$ is $\psi(\Phi_{\mu_j}, \epsilon)_{\omega_\theta} = \exp(\epsilon q(\theta, j))$. Then we have

$$\mathcal{E}_q^\epsilon(j|\theta) = \frac{\psi(\Phi_{\mu_j}, \epsilon)_{\omega_\theta} \frac{(\Psi'_{\mathcal{P},\epsilon} \Psi_{\mathcal{P},\epsilon})^{-1} \Psi'_{\mathcal{P},\epsilon} \mathbf{1})_j}{\mathbf{1}'(\Psi'_{\mathcal{P},\epsilon} \Psi_{\mathcal{P},\epsilon})^{-1} \Psi'_{\mathcal{P},\epsilon} \mathbf{1}}}{\sum_{k=1}^J \psi(\Phi_{\mu_k}, \epsilon)_{\omega_\theta} \frac{(\Psi'_{\mathcal{P},\epsilon} \Psi_{\mathcal{P},\epsilon})^{-1} \Psi'_{\mathcal{P},\epsilon} \mathbf{1})_k}{\mathbf{1}'(\Psi'_{\mathcal{P},\epsilon} \Psi_{\mathcal{P},\epsilon})^{-1} \Psi'_{\mathcal{P},\epsilon} \mathbf{1}}} = \frac{(\Psi_{\mathcal{P},\epsilon} \operatorname{diag}(\Psi'_{\mathcal{P},\epsilon} \Psi_{\mathcal{P},\epsilon})^{-1} \Psi'_{\mathcal{P},\epsilon} \mathbf{1}))_{\omega_\theta j}}{\sum_{k=1}^J (\Psi_{\mathcal{P},\epsilon} \operatorname{diag}(\Psi'_{\mathcal{P},\epsilon} \Psi_{\mathcal{P},\epsilon})^{-1} \Psi'_{\mathcal{P},\epsilon} \mathbf{1}))_{\omega_\theta k}}.$$

Letting $\sigma_{\tau^*,\epsilon}$ be as in Lemma 11, we have

$$(\Psi_{\mathcal{P},\epsilon}\operatorname{diag}(\Psi'_{\mathcal{P},\epsilon}\Psi_{\mathcal{P},\epsilon})^{-1}\Psi'_{\mathcal{P},\epsilon}\mathbf{1}))_{\omega j} = \sigma_{\tau^*,\epsilon}(j|\omega),$$
$$\Rightarrow \sum_{k=1}^{J}(\Psi_{\mathcal{P},\epsilon}\operatorname{diag}(\Psi'_{\mathcal{P},\epsilon}\Psi_{\mathcal{P},\epsilon})^{-1}\Psi'_{\mathcal{P},\epsilon}\mathbf{1}))_{\omega_\theta k} = \sum_{k=1}^{J}\sigma_{\tau^*,\epsilon}(k|\omega_\theta) = 1.$$

Then $\mathcal{E}^{\epsilon}_{q}(j|\theta) = \sigma_{\tau^*,\epsilon}(j|\omega_\theta)$ for each $j \in \{1,\dots,J\}$ and $\theta \in \{0,1\}^N$; the statement follows from Lemma 11. □

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 $\omega \in \Omega$,

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

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

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

The claim follows. □

Proof of Lemma 2 (Geometric Mechanisms are UPRR-Dominant) Suppose that $\mu \in \operatorname{supp}\tau^g_\epsilon$. 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 $\omega' \leq \omega'' \leq \omega^*(\mu)$,

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

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 $\omega'' \geq \omega' \geq \omega^*(\mu)$,

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

Fix $\mu' \in \operatorname{supp}\tau$. 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 $\omega' \leq \omega''$,

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

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, $\mu(\omega'')/\mu'(\omega'') \leq \mu(\omega')/\mu'(\omega')$. □

Proof of Lemma 3 (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

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

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

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

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

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

as desired.

(iii): Since A is compact and h is upper semicontinuous, $\max_{a \in A}(\sum_{w \in \Omega} \mu(w)h(w,a)) = \max_{a \in A} E_\mu[h(\omega,a)]$ exists for each $\mu \in \Delta(\Omega)$. By the law of iterated expectations, we have

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

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

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

as desired. □

Proof of Proposition 4 (UPRR-Dominance Implies Supermodular Stochastic Dominance) Let F and G denote the cumulative distribution functions of (W,X) and $(\hat{W},\hat{X})$, respec-

tively. Further, let

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

Since $\tau \succeq_{UPRR} \tau'$, 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 $\mu' \in \operatorname{supp} \tau'$, we have $\frac{\mu(\omega'')}{\mu(\omega')} \geq \frac{\mu'(\omega'')}{\mu'(\omega')}$ and $\frac{\mu(\tilde{\omega})}{\mu(\tilde{\omega})} \leq \frac{\mu'(\tilde{\tilde{\omega}})}{\mu'(\tilde{\omega})}$, or equivalently,

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

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 $\mu' \in \operatorname{supp} \tau'$,

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

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 $\mu' \in \operatorname{supp} \tau'$,

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

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

$$f(\omega'',y)g(\omega',y') - g(\omega'',y')f(\omega',y) = \sum_{\mu' \in t^{-1}(Q_{\tau',t}(y'))} \frac{\tau'(\mu)(f(\omega'',y)\mu'(\omega') - \mu'(\omega'')f(\omega',y))}{r_{\tau',t}(Q_{\tau',t}(y'))} \geq 0, \tag{17}$$

$$f(\tilde{\tilde{\omega}},z)g(\tilde{\omega},z') - g(\tilde{\tilde{\omega}},z')f(\tilde{\omega},z) = \sum_{\mu' \in t^{-1}(Q_{\tau',t}(z'))} \frac{\tau'(\mu)(f(\tilde{\tilde{\omega}},z)\mu'(\tilde{\omega}) - \mu'(\tilde{\tilde{\omega}})f(\tilde{\omega},z))}{r_{\tau',t}(Q_{\tau',t}(z'))} \leq 0. \tag{18}$$

Integrating (17) on $[x,1]\times[x,1]$ and (18) on $[0,x]\times[0,x]$ yields

$$\begin{aligned}\int_x^1 f(\omega'',y)dy\int_x^1 g(\omega',y')dy' &\geq \int_x^1 g(\omega'',y')dy'\int_x^1 f(\omega',y)dy\\ \Leftrightarrow \frac{\int_x^1 f(\omega'',y)dy}{\int_x^1 g(\omega'',y)dy} &\geq \frac{\int_x^1 f(\omega',y)dy}{\int_x^1 g(\omega',y)dy}; \qquad (19)\\ \int_0^x f(\tilde{\omega},z)dz\int_0^x g(\tilde{\omega},z')dz' &\leq \int_0^x g(\tilde{\omega},z')dz'\int_0^x f(\tilde{\omega},z)dz\\ \Leftrightarrow \frac{\int_0^x f(\tilde{\omega},z)dz}{\int_0^x g(\tilde{\omega},z)dz} &\leq \frac{\int_0^x f(\tilde{\omega},z)dz}{\int_0^x g(\tilde{\omega},z)dz}. \qquad (20)\end{aligned}$$

It follows from (20) 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 $\tilde{\tilde{\omega}}\in[\tilde{\omega},N]$,

$$\int_0^x g(\tilde{\omega},z)dz\geq\int_0^x f(\tilde{\omega},z)dz\Rightarrow\int_0^x g(\tilde{\tilde{\omega}},z)dz\geq\int_0^x f(\tilde{\tilde{\omega}},z)dz. \qquad (21)$$

Likewise, it follows from (19) 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 $\omega''\in[\omega',Q_{\tau,\omega^*}(x)]$,

$$\int_x^1 f(\omega',z)dz\geq\int_x^1 g(\omega',z)dz\Rightarrow\int_x^1 f(\omega'',z)dz\geq\int_x^1 g(\omega'',z)dz. \qquad (22)$$

By Lemma 3 (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 (22) 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 $\omega''\in[\omega',Q_{\tau,\omega^*}(x)]$,

$$\int_0^x g(\omega',z)dz\geq\int_0^x f(\omega',z)dz\Rightarrow\int_0^x g(\omega'',z)dz\geq\int_0^x f(\omega'',z)dz. \qquad (23)$$

Moreover, for each $x\in[0,1]$, $\tilde{\omega}\in[0,Q_{\tau,\omega^*}(x)]$, and $\tilde{\tilde{\omega}}\in[Q_{\tau,\omega^*}(x),N]$, (21) and (23) imply that

$$\begin{aligned}\int_0^x g(\tilde{\omega},z)dz\geq\int_0^x f(\tilde{\omega},z)dz &\Rightarrow\int_0^x g(Q_{\tau,\omega^*}(x),z)dz\geq\int_0^x f(Q_{\tau,\omega^*}(x),z)dz\\ &\Rightarrow\int_0^x g(\tilde{\tilde{\omega}},z)dz\geq\int_0^x f(\tilde{\tilde{\omega}},z)dz. \qquad (24)\end{aligned}$$

Then from (21), (23), and (24), 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: (21) holds for all $\tilde{\omega}, \tilde{\tilde{\omega}} \in \Omega$ with $\tilde{\omega} \leq \tilde{\tilde{\omega}}$. 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 3 (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 function $h : \Omega \times [0,1] \to \mathbb{R}$,

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

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 $\mu \in \operatorname{supp} \tau'$, 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 3 (ii),

$$E[h(\hat{W}, Q_{\tau',a^*}(\hat{X}))] = E_{\tau'}\left[E_\mu[h(\omega, a^*(\mu))]\right] = E_{\tau'}\left[\max_{a\in A} E_\mu[h(\omega,a)]\right]. \tag{25}$$

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

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

Let $\omega^* : \text{supp}\,\tau \to \mathbb{R}$ be the function mapping each posterior in supp τ to its peak used in the UPRR ordering $\tau \succeq_{UPRR} \tau'$, and let (W, X) be the Frechét representation of τ with respect to ω^*. Then by Proposition 4 and (25),

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

The statement then follows directly from Lemma 3 (iii). □

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