



HAL
open science

Universal optimality and robust utility bounds for metric differential privacy

Natasha Fernandes, Annabelle Mciver, Catuscia Palamidessi, Ming Ding

► **To cite this version:**

Natasha Fernandes, Annabelle Mciver, Catuscia Palamidessi, Ming Ding. Universal optimality and robust utility bounds for metric differential privacy. *Journal of Computer Security*, 2023, pp.1-42. 10.3233/JCS-230036 . hal-04349262

HAL Id: hal-04349262

<https://inria.hal.science/hal-04349262>

Submitted on 17 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Universal Optimality and Robust Utility Bounds for Metric Differential Privacy

Natasha Fernandes^{1,2,3}, Annabelle McIver¹, Catuscia Palamidessi², and Ming Ding³

¹Macquarie University, Sydney, Australia

²Institut Polytechnique de Paris, France

³Data61, Sydney, Australia

Abstract—We study the privacy-utility trade-off in the context of metric differential privacy. Ghosh et al. introduced the idea of *universal optimality* to characterise the “best” mechanism for a certain query that simultaneously satisfies (a fixed) ε -differential privacy constraint whilst at the same time providing better utility compared to any other ε -differentially private mechanism for the same query. They showed that the Geometric mechanism is “universally optimal” for the class of “counting queries”. On the other hand, Brenner and Nissim showed that outside the space of counting queries, and for the specific loss function called ℓ_{bin} , no such universally optimal mechanisms exist. Except for universal optimality of the Laplace mechanism, there have been no generalisations of these universally optimal results to other classes of differentially-private mechanisms.

In this paper we use metric differential privacy and quantitative information flow as the fundamental principle for studying universal optimality. Metric differential privacy is a generalisation of both standard (i.e., central) differential privacy and local differential privacy, and it is increasingly being used in various application domains, for instance in location privacy and in privacy preserving machine learning. As do Ghosh et al. and Brenner and Nissim, we measure utility in terms of loss functions, and we interpret the notion of a privacy mechanism as an information-theoretic channel satisfying constraints defined by ε -differential privacy and a metric meaningful to the underlying state space. Using this framework we are able to clarify Nissim and Brenner’s negative results, showing (a) that in fact *all* privacy types contain optimal mechanisms relative to certain kinds of non-trivial loss functions, and (b) extending and generalising their negative results beyond ℓ_{bin} specifically to a wide class of non-trivial loss functions.

Our exploration suggests that universally optimal mechanisms are indeed rare within privacy types. We therefore propose weaker universal benchmarks of utility called *privacy type capacities*. We show that such capacities always exist and can be computed using a convex optimisation algorithm.

We illustrate these ideas on a selection of examples with several different underlying metrics.

I. INTRODUCTION

One of the main challenges in privacy is to devise mechanisms that protect the sensitive information contained in the data, while at the same time preserving an acceptable degree of utility. Differential privacy [1], [2], one of the most successful frameworks for privacy protection, allows access to a dataset only via querying the dataset through an interface controlled by the data curator, which ensures privacy by adding controlled random noise to the result of the query before reporting it. In this setting, the trade-off with utility is determined by how

the noise is calibrated: a “bad” calibration may cause a lot of utility loss without necessarily improving privacy. Of course, the ideal mechanisms are the Pareto optimal ones, i.e., those that offer the best utility for a given level of privacy.

There are many notion of utility, depending on the goals and the means of the data consumers. The goals are usually formalised in terms of *loss functions*, and an important class of consumers is the so-called *Bayesian consumers*, who may have some prior knowledge about the data and are able to make the most out of the reported answer, in the sense of deriving (via Bayesian inference, using their prior) the information that minimises the loss with respect to the true answer.

In the context of the Bayesian notion of utility, an important desideratum for a mechanism is to be optimal for all consumers, no matter their prior. This is in line with the property of differential privacy, which does not depend on the prior knowledge of the attacker. This independence from the prior is appealing because when we design mechanisms we don’t know what kind of knowledge attackers and consumers may acquire over time.

The goal of this paper is study universally optimal mechanisms for Bayesian consumers over any class of loss functions (ie. not restricted to the monotone class). While monotone loss functions capture many consumers of interest, other research [3] provides examples of non-monotonic consumers which we argue should be included in a complete study of optimality. Moreover, the non-monotone ℓ_{nb} loss function (introduced later in this paper) is crucial in computing capacities for classes of privacy mechanisms.

To better understand the issues that arise in this context, let us introduce some notation.

Let \mathcal{C} be a class of (Bayesian) consumers, where each consumer is identified by a prior π and a loss function ℓ , and let \mathcal{D} be a class of ε -differentially private mechanisms. A mechanism $M \in \mathcal{D}$ probabilistically transforms the true answer (input) into a reported answer (output, observation)¹, and it is called *universally optimal* w.r.t. \mathcal{C} if M provides the best utility for every consumer in \mathcal{C} compared with any other $M' \in \mathcal{D}$. “Best utility” for a consumer w.r.t. a mechanism M

¹This functionality of the mechanism corresponds to the so-called “oblivious” model, in which the noise depends only on the result of the query and not on the original dataset.

is the minimal value of the consumer’s loss function over all possible remappings from observations to inputs, i.e.,

$$U_\ell(\pi, M) = \min_r \sum_{x: X} \pi_x \sum_{y: Y} M_{x,y} \ell(r(y), x) \quad (1)$$

for inputs X , observations Y and remapping function r . The universally optimal mechanism M is one for which $U_\ell(\pi, M) \leq U_\ell(\pi, M')$ for all other mechanisms $M' \in \mathcal{D}$, priors π and all loss functions ℓ within the class \mathcal{C} .

Ghosh et al. [4] found that a universally optimal mechanism exists for a particular class \mathcal{D} of “counting query” mechanisms and the class \mathcal{C} of “monotonic” consumers, i.e., described by a monotonic loss function. Conversely, Brenner and Nissim [5] showed that no such optimal mechanism exists for the same class of consumers with respect to various queries. Moreover, they proved that for the specific loss function ℓ_{bin} such optimal mechanisms never exist outside the linear structures of inputs corresponding to the counting queries.

The above results are in the context of the so-called *central model* of differential privacy, which assumes that the noise is added by a trusted curator. In recent years, however, the so-called *local model* [6] has become more and more popular, also due to the interest of large companies such as Google and Apple. In the local model, the noise is added directly by the data producer, so the microdata are already obfuscated. Consequently, this model is more robust w.r.t. potential security breaches, and there is no need for a trusted third party.

Recently it has been shown that both the central and local models can be unified under a generalised definition of differential privacy known as \mathbf{d} -privacy [7], which describes a *metric privacy type*. \mathbf{d} -privacy is increasingly being used in various application domains, for instance in location privacy, where it is known under the name of *geo-indistinguishability* [8], and in privacy preserving machine learning [9].

Definition 1 (Metric privacy type). *Given a metric space (X, \mathbf{d}) , we say that a mechanism $M : X \rightarrow \mathbb{D}\mathcal{Y}$ satisfies \mathbf{d} -privacy if for all $x, x' \in X$, the following constraint holds:*

$$M(x)(Y) \leq e^{\mathbf{d}(x,x')} M(x')(Y), \quad (2)$$

where we denote $M(x)(Y)$ the probability that M ’s output is contained in the subset $Y \subseteq \mathcal{Y}$. We call the class of \mathbf{d} -private mechanisms on X the metric privacy type associated to (X, \mathbf{d}) which we denote by $\mathcal{T}_{X, \mathbf{d}}$.

It is easy to see that \mathbf{d} -privacy subsumes both the central and the local model. For example, the metric privacy type for “counting queries” (in the central model) is $\mathcal{T}_{\mathbb{N}, \mathbf{d}_2}$ where \mathbf{d}_2 is the Euclidean distance and \mathbb{N} the natural numbers². For local differential privacy it is $\mathcal{T}_{X, \mathbf{d}_D}$ for any inputs X , where \mathbf{d}_D is the Discrete metric assigning distance 1 to all $x \neq x' \in X$. Also note that \mathbf{d} -privacy is strictly more general than local differential privacy, as it is able to express the distance between inputs and tune the noise accordingly. This is important, for

² M can also be seen as a mechanism from datasets to reported answers, as is customary in central differential privacy. In this case, \mathbf{d} would be the Hamming distance.

instance, in machine learning applications, where \mathbf{d} -privacy can privatise the gradient descent step without the need for clipping.

A. Contributions

Our contributions are as follows:

- 1) Using principles from Quantitative Information Flow, we study optimality for metric differential privacy in full generality; that is, \mathbf{d} -private mechanisms which provide the best utility over arbitrary classes of loss functions. This extends the study by Ghosh et al. [4] and Brenner and Nissim [5] which considered optimality only for the central model and only for the class of monotonic loss functions.
- 2) We provide a complete characterisation of optimality within a privacy type for finite X . This characterisation is independent of the workflow (oblivious, local or otherwise) in which the privacy type arises. This result is significant in that it provides a finite class of mechanisms which characterise the space.
- 3) We show that every privacy type contains universally optimal mechanisms (for non-trivial loss functions). This indicates that the impossibility result of Brenner and Nissim is too strong; we provide a tightening of their result, clarifying that the impossibility result for “sum queries” holds for a class of *strictly* monotonic consumers.
- 4) We extend the qualitative study of optimality to a quantitative study, finding robust capacity bounds for universally optimal mechanisms within a privacy type. These capacity results provide a tight quantitative bound on the maximum leakage – a measure of the usefulness of a mechanism within a privacy type. In practice for example, such bounds provide benchmarks for experimental evaluation of implemented mechanisms.

B. Related Work

The problem of universal optimality for arbitrary privacy types, as modelled by \mathbf{d} -privacy, has not yet been studied in full generality. Chatzikokolakis et al. [7] showed that optimal mechanisms for “sum queries” can be constructed in the oblivious setting by considering a Manhattan metric, rather than a Hamming metric, on the original datasets. Recent work by Fernandes et al. [10] considers universal optimality of the Laplace mechanism in the continuous setting, using the same QIF framework as in our paper, although our work focuses on a more general problem of universal optimality restricted to discrete spaces. Other optimality results have been found in non-Bayesian settings. Gupte et al. [11] proved a universal optimality result for a type of risk-averse consumer using a *minimax* formulation. Aharya et al. [12] prove an optimality result for binary mechanisms similar to our result of Cor. 9, but the focus of their work is on minimax consumers. Kairouz et al. [13] proved optimality for “staircase mechanisms” using f -divergence as a utility measure, while Koufogiannis et al. [14] showed that the Laplace mechanism is optimal for the

mean-squared error. Finally, Asi and Duchi [15] studied near-instance optimality, designed for particular dataset instances, in contrast with our study of optimality over mechanisms. It is important to note that the notions of utility induced by the above non-Bayesian consumers are simpler than ours, and their trade-off with privacy is more direct. Also, the question of universality w.r.t. the prior does not arise. Furthermore, most of the results assume heavy restrictions on the noise functions.

C. Organisation of the Paper

We begin in §II with a summary of the results needed from Quantitative Information Flow including how to formulate the Universal optimality problem under the same assumption of “Bayesian consumers” as used by Ghosh et al. and Brenner and Nissim. In §III we introduce privacy types based on a metric \mathbf{d} and a set of secrets \mathcal{X} ; given these ingredients we provide a complete characterisation of the mechanisms of that type. In §IV we provide a thorough investigation of utility in terms of loss functions and how there is a close relationship between the complexity of the loss function and the structure of the underlying privacy mechanisms from which we can deduce both when universal properties exist and when they do not depending on the loss function (eg. its complexity) and the composition of the privacy mechanisms. In §V we introduce privacy type capacities and show that they provide a benchmark for computing how much useful information can be extracted from a privacy mechanism. Finally in §VI we provide a range of examples that illustrate these ideas.

Where possible we have sketched proofs in the main body of the paper. Full proofs can be found in the appendix.

II. QUANTITATIVE INFORMATION FLOW FOR PRIVACY MECHANISMS

A probabilistic channel C takes an input $x \in \mathcal{X}$ and outputs an “observation” $y \in \mathcal{Y}$ according to a distribution $\mathbb{D}\mathcal{Y}$. In the discrete case, such channels are $\mathcal{X} \times \mathcal{Y}$ matrices C whose row- x , column- y element $C_{x,y}$ is the probability that input x will produce output y . The x -th row $C_{x,-}$ is thus a discrete distribution in $\mathbb{D}\mathcal{Y}$. We write $\mathcal{X} \rightarrow \mathbb{D}\mathcal{Y}$ for the channel type.

Given a *prior* distribution $\pi: \mathbb{D}\mathcal{X}$ on \mathcal{X} , the channel C can be applied to π to create a joint distribution J in $\mathbb{D}(\mathcal{X} \times \mathcal{Y})$, written $\pi \triangleright C$ and where $J_{x,y} := \pi_x C_{x,y}$. For that J , the left-marginal $\sum_y J_{x,y}$ (denoted $J_{x,\Sigma}$) gives the prior π again, i.e. the probability that the input was x — thus $\pi_x = J_{x,\Sigma}$. The right marginal $J_{\Sigma,y}$ is the probability that the output is y , given both π and C . The y -*posterior* distribution on \mathcal{X} is the *conditional* probability that the input was x if that y was output: it is the y -th column (of the joint J) divided by the marginal probability of that y , that is $J_{-,y}/J_{\Sigma,y}$ (provided the marginal is not zero).

If we fix π and C , and use the conventional abbreviation p_{XY} for the resulting joint distribution ($\pi \triangleright C$), then the usual notations for the above are p_X for left marginal ($=\pi$) and $p_X(x)$ for its value π_x at a particular x , with p_Y and $p_Y(y)$ similarly for the right marginal. Then $p_{X|y}(x)$ is the posterior probability of the original observation’s being x when y has

C	y_0	y_1	y_2
x_0	$\frac{2}{3}$	$\frac{1}{6}$	$\frac{1}{6}$
x_1	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$
x_2	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{2}{3}$

$$\Delta_C = \begin{bmatrix} 4/7 & 1/4 & 1/7 \\ 2/7 & 1/2 & 2/7 \\ 1/7 & 1/4 & 4/7 \\ 7/18 & 2/9 & 7/18 \end{bmatrix}$$

Fig. 1: A channel C and its corresponding hyper Δ_C produced by the action of the uniform distribution.

been observed. Further, we can write just $p(x)$ and $p(y)$ and $p(x|y)$ when context makes the (missing) subscripts clear.

A. Bayesian Adversaries and Generalised Entropy

QIF assumes that adversaries are Bayesian: they are equipped with a prior $\pi: \mathbb{D}\mathcal{X}$ over secrets \mathcal{X} and use their knowledge of the channel $C: \mathcal{X} \rightarrow \mathbb{D}\mathcal{Y}$ to maximise their advantage after making an observation. This is modelled in full generality using the “ g -leakage framework” [16]. In detail, a gain function $g: \mathcal{W} \times \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ models the gain to an adversary who takes an *action* $w \in \mathcal{W}$ when the value of the secret is $x \in \mathcal{X}$. Focussing on actions rather than observations naturally encompasses the idea of “remapping” observations to their associated most likely secret value. Moreover g -leakage gives a measurement for how successful are (privacy) mechanisms at communicating useful information by comparing the “prior” and “posterior” g -vulnerabilities as follows. For gain function g the (prior) vulnerability wrt. the prior π (for a consumer/adversary with this information) is the maximum gain over all possible actions: $V_g(\pi) := \max_{w: \mathcal{W}} \sum_{x: \mathcal{X}} \pi_x g(w, x)$; it quantifies the adversary’s success in the scenario defined by π and g . The expected (*posterior*) *vulnerability* is the adversary’s expected gain after observing y is $V_g[\pi \triangleright C] := \sum_{y: \mathcal{Y}} p(y) V_g(p(\mathcal{X}|y))$, equivalently:

$$V_g[\pi \triangleright C] = \sum_{y: \mathcal{Y}} \max_{w: \mathcal{W}} \sum_{x: \mathcal{X}} \pi_x C_{x,y} g(w, x). \quad (3)$$

The greater the difference in the prior/posterior vulnerability, the better is the adversary able to use the transmitted to infer the value of the secret. (See §II-B below.)

Note that we also use an equivalent formulation of leakage in terms of loss functions and minimising “generalised entropies”. In particular a function $\ell: \mathcal{W} \times \mathcal{X} \rightarrow \mathbb{R}$ (also) defines a generalised entropy $U_\ell: \mathbb{D}\mathcal{X} \rightarrow \mathbb{R}$:

$$U_\ell(\pi) := \min_{w: \mathcal{W}} \sum_{x: \mathcal{X}} \pi_x \ell(w, x), \quad (4)$$

with the formulation for expected loss through a channel defined similarly using minimisation (compare (3)):

$$U_\ell[\pi \triangleright C] = \sum_{y: \mathcal{Y}} \min_{w: \mathcal{W}} \sum_{x: \mathcal{X}} \pi_x C_{x,y} \ell(w, x). \quad (5)$$

We will often use g for gain and ℓ for loss, but either can be used in formulating a vulnerability (V_g/V_ℓ) or generalised entropy (U_g/U_ℓ); the leakage theory based on losses or gains is equivalent [3].

Remark 1. We note that the Bayesian formulation of utility loss adopted in the literature ([4], [5]) and presented in Eqn 1, is equivalent to the formulation for posterior utility given in Eqn 5 (see Appendix §A for proof).

We adopt Eqn (5) to describe utility for consumers.

Equation (3) indicates that the actual output \mathcal{Y} can be abstracted, leaving the leakage properties of a channel in the context of a prior to be determined by a *hyper-distribution* (or simply “hyper”). A hyper is a distribution of distributions on \mathcal{X} , having type $\mathbb{D}\mathcal{D}\mathcal{X}$ or $\mathbb{D}^2\mathcal{X}$. Given a channel C and prior π , we write $[\pi \triangleright C]$ for the hyper whose support is the set of posterior distributions $p(\mathcal{X}|y)$ on \mathcal{X} and which assigns the corresponding marginal $p(y)$ to each. We usually denote by δ^y the posterior $p(\mathcal{X}|y)$ and by α_y its corresponding marginal.

An example of a channel and hyper is shown in Figure 1. The hyper Δ_C is produced by pushing the uniform prior through the channel C to produce a joint distribution, which is then marginalised along its columns to produce the posteriors δ^y (the columns in Δ_C) and the corresponding y -marginals α_y (the labels beneath each column). The prior π can be recovered by averaging the posteriors δ_y via the marginals α_y , revealing a significant correspondence between channels and hypers:

Lemma 1 (Cor. 4.8 from [3]). *For any channel C and prior π , $\pi = \sum_i \alpha_i \delta^i$ where $[\pi \triangleright C] = \sum_i \alpha_i [\delta^i]$.*

Figure 1 illustrates this relationship between a channel and its associated hyper.

It turns out that hypers, and their refinement relation, have a compelling *geometric* interpretation (first explained in [3, Ch. 12]), which we use below to study optimal utility.

B. Robustness measures: refinement and channel capacities

The study of g -vulnerabilities has given rise to an elegant theory of *refinement*: we say that channel A refines channel B , written $B \sqsubseteq A$, to mean that A is safer than B which holds exactly when the g -vulnerability of A is no greater than that of B for *any* prior or gain function, i.e.,

$$B \sqsubseteq A \text{ iff } V_g[\pi \triangleright A] \leq V_g[\pi \triangleright B] \text{ for all } \pi: \mathbb{D}\mathcal{X}, g \in \mathbb{G}\mathcal{X} \quad (6)$$

The above holds dually for $U_\ell[\pi \triangleright A] \geq U_\ell[\pi \triangleright B]$. Remarkably, this holds whenever there exists a channel P such that $BP = A$ (writing BP for matrix multiplication) [17]. In fact P corresponds to a “postprocessing” of the output of a channel, always suppressing information flow. The channel \mathbb{I} with a single column of 1’s is maximal in the refinement order, and leaks nothing at all i.e. $U_\ell[\pi \triangleright \mathbb{I}] = U_\ell(\pi)$.

Note that refinement is a strong universal utility property of channels (and therefore mechanisms modelled as channels). Refinement between two channels A and B means that for all priors and all gain (loss) functions channel A will always have more gain (less loss) when compared to the corresponding scenario for channel B . We shall see therefore that when A, B are members of the same privacy type then a refinement relation between them implies a universal optimality property between the two. We find, however, that in general there is no

mechanism that anti-refines all mechanisms within the type, and therefore the universal optimality property can only hold for a subset of specific gain/loss functions.

Equivalently, refinement can be expressed as a relation between hypers: $[\pi \triangleright B] \sqsubseteq [\pi \triangleright A]$ exactly when there is a “refining Earth Move” (defined below) from the posteriors of $[\pi \triangleright B]$ to the posteriors of $[\pi \triangleright A]$. We recall the important correspondence between refinement of channels and refinement of hypers [3, Ch. 12]:

$$B \sqsubseteq A \text{ iff } [\pi \triangleright B] \sqsubseteq [\pi \triangleright A] \quad (7)$$

where π is *any* full support prior. In other words, refinement of channels (on all priors) is characterised by refinement of hypers on a single (full support) prior.

In this paper we use this correspondence to characterise the space of \mathbf{d} -private mechanisms and those which are universally ℓ -optimal. We write Δ etc. for general hypers on $\mathbb{D}^2\mathcal{X}$; we note that we can depict Δ as a weighted sum of vectors, where we identify each posterior δ^y in the support of Δ as a point in $\mathbb{R}^{|\mathcal{X}|}$. We make the following observations [3, Ch. 12]):

- A “refining Earth Move” [18] from Δ_B to Δ_A exists when each posterior δ^y of Δ_A is realised as a convex combination of posteriors of Δ_B (ie. an “Earth Move” from Δ_B to each δ^y) which exactly match the marginals of Δ_B to produce the corresponding marginals on Δ_A .
- If Δ_B ’s posteriors are linearly independent (as vectors), the Earth Move from a) exists whenever the posteriors of Δ_A lie inside the convex hull of the posteriors of Δ_B .
- The difference between the posterior and prior vulnerabilities gives a measure of the accuracy of the channel. The additive g -leakage is $|V_g[\pi] - V_g[\pi \triangleright A]|$, and the multiplicative g -leakage is $V_g[\pi \triangleright A]/V_g[\pi]$. The corresponding channel capacities take the maximal difference over all priors and $\mathbb{G}\mathcal{X}$, the 0/1-bounded vulnerabilities.³ Specifically the multiplicative and additive capacities respectively are:

$$\mathcal{ML}^\times(C) := \max_{\pi, g} (V_g[\pi \triangleright C]/V_g[\pi]) \quad (8)$$

and

$$\mathcal{ML}^+(C) := \max_{\pi, g} |V_g[\pi] - V_g[\pi \triangleright C]| \quad (9)$$

It is easy to see that $1 \leq \mathcal{ML}^\times(C) \leq |\mathcal{X}|$ and $0 \leq \mathcal{ML}^+(C) \leq 1$. The larger either capacity, the more accurate the channel. For example, $\mathcal{ML}^+(\mathbb{I}) = 0$, and $\mathcal{ML}^\times(\mathbb{I}) = 1$ consistent with it transmitting no information at all. Capacities provide an alternative robust bound on the ability of any mechanism in a privacy type to deliver accurate information since it represents a tight upper bound on the maximum leakage of any channel in the type. We show how to compute capacities for the whole type either by proving a universal optimality result or using the characterisation given in §III below.

We shall show that irrespective of whether a type can support universal optimality results, the capacities are well-defined and there exist mechanisms that achieve the capacity.

³i.e. $0 \leq V_g[\pi] \leq 1$

Such mechanisms represent the most efficient implementations that respect the privacy threshold whilst at the same time delivering the most information about the utility.

C. Universal optimality and refinement

Given the equivalence between utility notions in the literature and the QIF notion of posterior ℓ -uncertainty (Remark 1), there is then a natural connection between utility and refinement which extends to universal optimality. In its most general form, we define universal optimality as:

Definition 2. (*Universal optimality*) Given a class of mechanisms \mathcal{T} we say that $M \in \mathcal{T}$ is universally optimal (wrt. \mathcal{T}) if $U_\ell[\pi \triangleright M] \leq U_\ell[\pi \triangleright M']$ for all $M' \in \mathcal{T}$, all loss functions ℓ and all priors $\pi \in \mathbb{D}\mathcal{X}$.

Using Eqn 6, this says that $M \in \mathcal{T}$ is universally optimal whenever $M' \sqsubseteq M$ for all $M' \in \mathcal{T}$. This definition appears to be too strong; indeed most mechanisms –even within the same privacy type– are not related by refinement [19]. Surprisingly, we find that there is a class \mathcal{T} for which this strong notion of universal optimality holds (Cor. 9). However, we also find that this class is unique (Thm. 11) and universal optimality is too strong in general.

Therefore, following Ghosh et al. we study a weaker order defined by specific loss functions:

Definition 3. (*Universal ℓ -optimality*) Given a loss function ℓ and a class of mechanisms \mathcal{T} we say that $M \in \mathcal{T}$ is universally ℓ -optimal (wrt. \mathcal{T}) if $U_\ell[\pi \triangleright M] \leq U_\ell[\pi \triangleright M']$ for all $M' \in \mathcal{T}$ and priors $\pi \in \mathbb{D}\mathcal{X}$.

The universal optimality definition adopted by Ghosh et al. is, using Def. 3, universal ℓ -optimality for all “monotone” loss functions ℓ (see Def. 6 ahead). We can think of Def. 3 as describing a type of restricted refinement which holds for all priors but only for loss functions $\ell \in \mathcal{L}$. This idea of refinement restricted to loss function classes has not to our knowledge been previously studied, and is the focus of §IV.

III. CHARACTERISING METRIC PRIVACY TYPES

In this section we study mechanisms within a privacy type, finding a new characterisation of the type requiring only a finite set of distinguished mechanisms which we call *kernel mechanisms*. This characterisation captures the leakage properties of the privacy type, reducing the study of optimality to the study of kernel mechanisms.

A. Privacy types as hyper-distributions

Following [20], we write a mechanism M of privacy type $\mathcal{T}_{\mathcal{X}, \mathbf{d}}$ as a channel $C: \mathcal{X} \rightarrow \mathbb{D}\mathcal{Y}$ satisfying $C_{x,y} \leq e^{\mathbf{d}(x,x')} C_{x',y}$ for all $x, x' \in \mathcal{X}$ and $y \in \mathcal{Y}$. ie., As a set of constraints on each column of the channel. Importantly, there is a correspondence between C and hyper-distributions as follows:

Lemma 2. C is \mathbf{d} -private iff $\delta_x^y - e^{\mathbf{d}(x,x')} \delta_{x'}^y \leq 0$ for every δ^y in the support of $[v \triangleright C]$ where $v \in \mathbb{D}\mathcal{X}$ is the uniform distribution.

Applying Lem. 2 to each pair x, x' yields a feasible region of hypers (defined by their posteriors δ^y) enclosed by hyperplanes –therefore convex– which we call “the space of $\mathcal{T}_{\mathcal{X}, \mathbf{d}}$ hypers”. Figure 2 illustrates one such space.

Observe that although Lem. 2 specifies the uniform distribution, Eqn (7) says that this restriction does not apply to channel refinement. In other words, we can reason about channel refinement (over all priors) by reasoning in the space of hypers w.r.t. the uniform distribution. This differentiates our approach from standard convex optimisation formulations of the differential privacy constraints on channels [21] which require quantification over all priors.




Fig. 2: The space of $\mathcal{T}_{\mathcal{X}, \mathbf{d}_2}$ hypers on $\mathcal{X} = \{x_1, x_2, x_3\}$ for $\varepsilon = \ln 2$ and $\mathbf{d}_2(x_1, x_2) = 1$, $\mathbf{d}_2(x_2, x_3) = 1$, $\mathbf{d}_2(x_1, x_3) = 2$. The coloured hyperplanes are the linear constraints which intersect the (triangular) simplex in \mathbb{R}^3 , forming the convex region shown by the dotted lines. Every point inside this region represents a $\ln 2$ - \mathbf{d}_2 -private posterior, and any $\ln 2$ - \mathbf{d}_2 -mechanism (channel) corresponds to a convex sum of posteriors that average to the uniform prior v .

B. A complete characterisation of privacy types

The convex space above describes the complete set of potential posteriors which can be found in hyper-distributions $[v \triangleright C]$. Using Lem. 1, we can construct valid channels from a set of posteriors, provided that they contain the uniform distribution and can be averaged to the uniform prior. Moreover, Lem. 2 says that if the posteriors satisfy the \mathbf{d} -privacy constraints, then so will the corresponding channel. (Notice this in Figure 1: the ratio between elements in each column of C matches the corresponding ratios in Δ_C). This means we can construct \mathbf{d} -private channels from sets of \mathbf{d} -private posteriors. We use this fact now to examine optimality via hypers.

We distinguish 2 types of hypers (and their corresponding mechanisms) formed from vertices in the space of $\mathcal{T}_{\mathcal{X}, \mathbf{d}}$ hypers:

Definition 4. (*Vertex Mechanism/Hyper*) Let $V \in \mathcal{T}_{\mathcal{X}, \mathbf{d}}$ be a mechanism with corresponding hyper $\Delta = [v \triangleright V]$. We say that V is a vertex mechanism (and Δ a vertex hyper) if the inners in $[\Delta]$ are vertices in the space of $\mathcal{T}_{\mathcal{X}, \mathbf{d}}$ hypers.

Vertex mechanisms are of particular interest because of their refinement properties: they are the minimal elements in the refinement order in their privacy type (that is, they

have no anti-refinements which satisfy the privacy constraints). Put another way, it means that every $\mathcal{T}_{\mathcal{X},\mathbf{d}}$ mechanism is a refinement of a vertex mechanism. (See Appendix §B for details). However, (potentially) infinitely many exist in any privacy type. We will instead make use of particular vertex mechanisms that we call *kernel mechanisms*.

Definition 5. (*Kernel Mechanism/Hyper*) Let $K \in \mathcal{T}_{\mathcal{X},\mathbf{d}}$ have corresponding hyper $\Delta = [v \triangleright K]$. We say that K is a kernel mechanism and Δ a kernel hyper if K is a vertex mechanism, and if the inners in $[\Delta]$ are linearly independent (as vectors).

Importantly, every vertex mechanism can be expressed as a convex sum of kernel mechanisms, and these generate all $\mathcal{T}_{\mathcal{X},\mathbf{d}}$ mechanisms. This leads to the main result of this section.

Theorem 3. (*Fundamental Characterisation of $\mathcal{T}_{\mathcal{X},\mathbf{d}}$ Mechanisms*) Every $\mathcal{T}_{\mathcal{X},\mathbf{d}}$ mechanism is a refinement of a convex sum of kernel mechanisms.

The usefulness of this characterisation is that we can now explore questions about optimality using a finite set of kernel mechanisms as a basis for generating classes of universally optimal mechanisms.

Note that our notion of privacy types plays a role similar to that of “privacy constraint graph” in [5], but it is strictly more general. First of all, the “privacy (constraint) graph” is defined only for the central model, while, as explained in the introduction, the privacy types of \mathbf{d} -privacy subsume and generalise both the central and the local model. Furthermore, from a technical point of view, privacy graphs can be expressed using a metric, but the vice-versa is not true, since in privacy graphs there is no notion of distance between nodes.

IV. UTILITY FOR BAYESIAN CONSUMERS

As mentioned in §II we shall use loss functions \mathcal{L} (dual to gain functions) to determine the average accuracy (utility) of a mechanism with respect to a Bayesian consumer.

A. Classes of loss functions

A loss function defines a *generalised entropy* of type $\mathbb{D}\mathcal{X} \rightarrow \mathbb{R}$. Such entropies are concave (and continuous), i.e. $U_\ell(\pi_1 + p\pi_2) \geq pU_\ell(\pi_1) + (1-p)U_\ell(\pi_2)$. There are many loss functions that have been identified as useful for analysing information flow properties, including accuracy of mechanisms.

Definition 6 (Monotone, strictly monotone and trivial loss functions). *Loss functions are said to be monotone in \mathbf{d} denoted $\mathcal{L}_{\mathbf{d}}^m$, when they have the form $\ell(w, x) = m(\mathbf{d}(x, \alpha(w)))$, where m is a monotone, non-decreasing function on the reals, and $\alpha : \mathcal{W} \rightarrow \mathcal{W}$ is an injective function. The strict monotone subclass has m a strictly increasing function. The trivial loss functions \mathcal{L}^* are independent of \mathcal{W} , so have the form $f(x)$ for some real-valued function f . We have the following relationships between these loss functions:*

$$\mathcal{L}^* \subseteq \mathcal{L}_{\mathbf{d}}^{m^+} \subseteq \mathcal{L}_{\mathbf{d}}^m \subseteq \mathcal{L}.$$

Observe that the trivial loss functions are independent of the underlying metric and cannot be used to measure any

non-trivial information flow property because they correspond to adversarial settings where the adversary does not have a choice of actions. The associated entropies of trivial loss functions, namely U_ℓ correspond to linear functions in $\mathbb{D}\mathcal{X}$, i.e. $U_\ell(p \times \pi_1 + (1-p) \times \pi_2) = p \times U_\ell(\pi_1) + (1-p) \times U_\ell(\pi_2)$, and in fact $U_\ell[v \triangleright M] = U_\ell[v]$ whatever the mechanism M . Non-trivial loss functions –corresponding to non-planar generalised entropies– are critical for measuring information flow properties as they characterise how much an adversary is able to use partial information leaks to further his intent. For example, ℓ_{nib} (defined in Eqn (11)) can be seen to be non-trivial as depicted in Figure 3 at right: the planar regions represent the action taken by the adversary to *minimise* his loss relative to his (current) knowledge of the secret as represented by eg. a posterior distribution.

Ghosh et al. [4] were the first to identify some important structures in the relation between loss functions and privacy mechanisms. They identified the class of monotone loss functions relative to the privacy type $\overline{\mathcal{T}}_{\mathbb{N},\mathbf{d}_2}$. In their formulation the function α acts as a “remapping” to select the most likely input for a given observation made as part of the privacy mechanism (here represented as a channel).

Finally we note that we shall see examples later that show that the consideration of whether a loss function is included in the monotone set is sensitive to the underlying metric.

B. Examples of significant loss functions

The function ℓ_{bin} is a loss function commonly used to measure the ability of an adversary to guess the secret: note that its corresponding entropy $U_{\ell_{\text{bin}}}$ is referred to as *Bayes’ Risk*. Letting $\mathcal{W} = \mathcal{X}$, we define:

$$\ell_{\text{bin}}(x, w) := 0 \text{ if } (x = w) \text{ else } 1. \quad (10)$$

Its dual is defined

$$\ell_{\text{nib}}(x, w) := 1 \text{ if } (x = w) \text{ else } 0. \quad (11)$$

Observe that $\ell_{\text{bin}} \in \mathcal{L}_{\mathbf{d}_2}^m$, but $\ell_{\text{bin}} \notin \mathcal{L}_{\mathbf{d}_2}^{m^+}$, whilst $\ell_{\text{bin}} \in \mathcal{L}_{\mathbf{d}_D}^{m^+}$. Meanwhile ℓ_{nib} is not monotone in any metric, but it is non-trivial and we shall see that it is significant for computing channel capacities.

When \mathcal{W} and \mathcal{X} are subsets of a Euclidean space, we can define an average loss function:

$$\ell_{\text{Avg}}(x, w) := \mathbf{d}_2(w, x). \quad (12)$$

Note that ℓ_{Avg} is strictly monotone in the type $\overline{\mathcal{T}}_{\mathbb{N},\mathbf{d}_2}$; the corresponding entropy $U_{\ell_{\text{Avg}}}$ computes the average distance from the true value of the secret.

Finally we also consider general non-monotone loss functions, where ℓ is an arbitrary real-valued loss function. Although these have not yet been studied for privacy mechanisms, they offer significant insights because they can express generic properties for studying property inferences [3], [22], and they can provide quantitative robustness measures due to the “miracle theorems” which quantify channel capacity. We recall these important results as follows.




Fig. 3: $U_{\ell_{\text{nb}}}(\pi)$ for $X = \{x_1, x_2, x_3\}$. We can see that ℓ_{nb} is non-trivial because $U_{\ell_{\text{nb}}}$ is not planar.

Theorem 4 (Multiplicative Miracle theorem [3]). *The following⁴ holds for all mechanisms M , non-negative gain functions g and priors π :*

$$V_g[\pi \triangleright M] / V_g[\pi] \leq V_{\ell_{\text{nb}}}[\nu \triangleright M] / V_{\ell_{\text{nb}}}[\nu],$$

where $V_{\ell_{\text{nb}}}[\nu \triangleright M] = 1 - U_{\ell_{\text{nb}}}[\nu \triangleright M]$.

Theorem 5 (Additive Miracle theorem [3]). *The following holds for all mechanisms M , 0/1-bounded vulnerabilities V_g and priors π :*

$$V_g[\pi \triangleright M] - V_g[\pi] \leq 1 - |\mathcal{X}| U_{\ell_{\text{nb}}}[\nu \triangleright M].$$

Observe that Thm. 4 and Thm. 5 can give robust *quantitative bounds* on general information leakage properties (including utility) for privacy types discussed above. We will illustrate this for the Euclidean and Discrete metrics below.

C. Which mechanisms have non-trivial ℓ -optimal loss functions?

One way to understand the ℓ -optimality problem within a type is to make the association of the loss function with an adversary (or consumer as in Ghosh et al. [4]) explicit. Recall that the loss function formalises an adversary's intent by describing the cost of taking action w when the secret is x . This means that a loss function is actually universally optimal for mechanism M exactly when the adversary (formalised by ℓ) can best make use of the information leaked by M compared with any other mechanism in the privacy type. These adversaries prefer M over all other mechanisms in the privacy type, independent of their prior knowledge. Thus rather than starting with a loss function ℓ and asking whether it has an ℓ -optimal mechanism, instead given a mechanism M ask which adversaries (loss functions) would prefer M over all others. We study that idea next.

Definition 7 (Utility set). *Let $M \in \mathcal{T}_{X, \mathbf{d}}$. Define $L_M^{\mathcal{T}_{X, \mathbf{d}}} := \{\ell \in \mathcal{L} \mid M \text{ is universally } \ell\text{-optimal within } \mathcal{T}_{X, \mathbf{d}}\}$*

⁴Note that multiplicative capacity is defined by maximising, hence V_g .

Notice that $L_M^{\mathcal{T}_{X, \mathbf{d}}}$ is closed under addition and multiplication by a non-negative generalised scalar, where addition is defined:

$$(\ell_1 + \ell_2)(x, (w_1, w_2)) := \ell_1(x, w_1) + \ell_2(x, w_2), \quad (13)$$

and generalised scalar multiplication is defined for (non-negative) real-valued function v :

$$(v \times \ell)(x, w) := v(x) \times \ell(x, w). \quad (14)$$

Observe also that $L_M^{\mathcal{T}_{X, \mathbf{d}}}$ always contains \mathcal{L}^* the set of “trivial” loss functions, and the utility set for the trivial mechanism that leaks no information at all consists only of trivial loss functions.

More interestingly it turns out that the structure of a privacy mechanism is reflected in the structure of the corresponding utility sets. There are two important constructors for mechanisms where we see this relationship clearly. The first is the *external probabilistic choice*: for channels $M, M' \in \mathcal{T}_{X, \mathbf{d}}$ we write $M_p \oplus M'$ as the mechanism obtained by applying M with probability p and M' with probability $(1-p)$. This yields a channel containing the columns of M scaled by p together with the columns of M' scaled by $(1-p)$, whose columns therefore clearly satisfy the same differentially private constraints as M, M' and so $M_p \oplus M' \in \mathcal{T}_{X, \mathbf{d}}$.

The second is *restriction* of a mechanism in $\mathcal{T}_{X, \mathbf{d}}$ to one in $\mathcal{T}_{X', \mathbf{d}}$ where $X' \subset X$. We write $M \downarrow X'$ for the restriction of the mechanism $M \in \mathcal{T}_{X, \mathbf{d}}$ to operate only on the subset of secrets X' obtained by removing rows of M corresponding to the secrets $X - X'$. Notice that $M \downarrow X'$ retains the original metric \mathbf{d} because the differentially-private constraints are only evaluated on the rows corresponding to secrets in X' .

Correspondingly in the space of loss functions we have the following constructions. Denote by $\ell \downarrow X'$ the restriction of the loss function ℓ to some subset X' of its secrets, defined by $(\ell \downarrow X')(w, x) := \ell(w, x)$ for $x \in X'$. Conversely we can *lift* a loss function to a superset of secrets X , denoted $\ell \uparrow X$; it is obtained by “padding” the loss function with 0's. That is, $(\ell \uparrow X)(w, x) := \ell(w, x)$ when $x \in X \subseteq \mathcal{X}$ and 0 when $x \notin X$.

We observe now that utility sets have a structure that is dual to the structure of mechanisms.

Theorem 6 (ℓ -Optimal Duality). *Let $\mathcal{T}_{X, \mathbf{d}}$ and $\mathcal{T}_{X', \mathbf{d}}$ be privacy types where $X' \subseteq X$. For mechanisms $M, M' \in \mathcal{T}_{X, \mathbf{d}}$, we have the following.*

- 1) $\mathcal{L}^* \subseteq L_M^{\mathcal{T}_{X, \mathbf{d}}}$,
- 2) $M \sqsubseteq M'$ implies $L_M^{\mathcal{T}_{X, \mathbf{d}}} \supseteq L_{M'}^{\mathcal{T}_{X, \mathbf{d}}}$,
- 3) $L_{\mathbb{I}}^{\mathcal{T}_{X, \mathbf{d}}} = \mathcal{L}^*$, where \mathbb{I} is the trivial mechanism,
- 4) $L_{M \oplus_p M'}^{\mathcal{T}_{X, \mathbf{d}}} = L_M^{\mathcal{T}_{X, \mathbf{d}}} \cap L_{M'}^{\mathcal{T}_{X, \mathbf{d}}}$ for $0 < p < 1$,
- 5) $\ell \in L_{M \downarrow X'}^{\mathcal{T}_{X, \mathbf{d}}}$ if and only if $\ell \uparrow X \in L_M^{\mathcal{T}_{X, \mathbf{d}}}$.

Notice that Thm. 6(1) means that $L_M^{\mathcal{T}_{X, \mathbf{d}}}$ is always non-empty, containing at least the set of trivial loss functions. As might be expected, Thm. 6(2) implies that the more information leaked through a mechanism, the more loss functions it has in its utility set, with Thm. 6(3) showing that the trivial mechanism, which always implements differential privacy (since it leaks

no information at all), is only universally ℓ -optimal for the set of trivial loss functions.

Next, Thm. 6(4) shows that when mechanisms are constructed using external probabilistic choice, the resulting utility set is the intersection of the utility sets of the components.

Finally, Thm. 6(5) provides an important method for discovering ℓ -optimal results by studying simpler privacy types with fewer secrets, but having the same underlying metric. We make use of Thm. 6(5) in §VI when we find universal optimality results in various metric spaces of interest.

A corollary to Thm. 6 is that the kernel mechanisms generate maximal utility sets.

Corollary 7. *If $M \in \mathcal{T}_{X,d}$ is a universally ℓ -optimal mechanism then $\ell \in L_K^{\mathcal{T}_{X,d}}$ for some kernel mechanism K in the type.*

Proof: Follows from Thm. 6(2) and (4) as Thm. 3 implies M refines a convex sum of kernels. ■

But even more, in cases where there is a unique kernel mechanism generating the privacy type, that mechanism satisfies “universal optimality”.

Theorem 8. *If $K \in \mathcal{T}_{X,d}$ is the unique kernel mechanism then $L_K^{\mathcal{T}_{X,d}} = \mathcal{L}$.*

Proof: From Thm. 3 we note that all mechanisms $M \in \mathcal{T}_{X,d}$ must satisfy $K \sqsubseteq M$ and so by the definition of \sqsubseteq (e.g. (6) for loss functions) the result follows. ■

It turns out that there is exactly one space for which a unique kernel mechanism exists, and from Thm. 8 this yields the following universal optimality result:

Corollary 9. *If $|X| = 2$, and for any metric \mathbf{d} , the following mechanism (represented as a channel) is universally optimal:*

$$T = \begin{bmatrix} k \cdot e^{\mathbf{d}} & k \\ k & k \cdot e^{\mathbf{d}} \end{bmatrix}, \quad (15)$$

where $k = 1/(1+e^{\mathbf{d}})$ is a scaling factor to ensure rows sum to 1.

Proof: The space of $\mathcal{T}_{X,d}$ hypers on 2 inputs $\{x_1, x_2\}$ is an interval on the line $x_1 + x_2 = 1$ with only 2 vertices, corresponding to $x_1 = e^{\mathbf{d}}x_2$ and $x_2 = e^{\mathbf{d}}x_1$. Since these are linearly independent, they must be the posteriors of a kernel hyper (cf. Def. 5) which is the only vertex hyper in the space. The result follows from Thm. 8. ■

The above result is significant in that it holds for all loss functions, and is non-trivial in the sense that binary mechanisms are of interest in the privacy community. (eg., the original Random Response mechanism of Warner which is used in Google’s RAPPOR [23]).

We shall see that using *restriction* and Thm. 6 we will be able to build non-trivial loss functions contained in the utility sets for kernel mechanisms for any privacy type.

D. Robustness results

As mentioned earlier, Brenner and Nissim [5] studied the question of universal ℓ -optimality specifically for types described in terms of “privacy constraint graphs”, which for us

can be described by using metrics. For example their “sum queries” can be considered as mechanisms in the privacy type \mathcal{T}_{X,d_D} . Whilst their negative results apply only to ℓ_{bin} within their characterisation of privacy types in terms of constraint graphs, they leave open the more fundamental question about whether optimality for non-trivial loss functions exists more generally. The next theorem (partially) answers that question, showing that non-trivial ℓ -optimal mechanisms always exist even when \mathbf{d} is the Discrete metric, and as a consequence universal ℓ -optimality is more common than expected.

Theorem 10 (Universal optimality existence). *Given any privacy type $\mathcal{T}_{X,d}$, if $|X| > 1$ then there is a non-trivial monotone loss function ℓ and a mechanism $M \in \mathcal{T}_{X,d}$ such that M is universally ℓ -optimal.*

In particular Thm. 10 applies even when the underlying metric is \mathbf{d}_D . More significantly, once we have found basic non-trivial loss functions contained in a mechanisms utility set, we can construct more complex loss functions, also in the utility set, by using addition (Eqn (13)) and scaling (Eqn (14)). Moreover, depending on the underlying metric many kinds of loss function can be significant for evaluating accuracy. In fact ℓ_{bin} and ℓ_{nib} are particularly significant since, if a corresponding universally optimal mechanism exists within a type, a robust benchmark for optimal loss exists (see Thm. 15).

E. Impossibility results

The universal impossibility results of Brenner and Nissim were studied for the monotone loss function class wrt “sum queries”. In this section we generalise this result to arbitrary metrics and loss functions. Importantly, we tighten the impossibility, showing that it holds only for the class of *strictly* monotone loss functions. We note that our proof is more general and direct than Brenner and Nissim’s result as we are able to access generic properties of the QIF framework.

We begin with a fundamental impossibility result:

Theorem 11 (Impossibility of Universally Optimal Mechanisms). *There are no universally optimal mechanisms in $\mathcal{T}_{X,d}$ for any \mathbf{d} and for $|X| > 2$.*

Thm. 11 applies to all loss function classes, and thus is a strengthening of the Brenner and Nissim result in that it shows that the impossibility also applies to the “counting query” class of Ghosh et al.

In terms of the weaker universal ℓ -optimality, Brenner and Nissim’s impossibility result for the Bayes’ risk loss function ℓ_{bin} shows that, at least in the Discrete domain, there exist loss functions (and therefore consumers) for which no mechanisms are universally optimal. However, as we have seen, this is not the case for all loss functions, nor is it the case for all monotone loss functions. The next result tightens the impossibility for sum queries, showing that it holds for the class of strictly monotone loss functions.

Theorem 12. *Let ℓ be a strictly monotone loss function on \mathbf{d} for some metric \mathbf{d} . Then there are no universally ℓ -optimal mechanisms in the privacy type \mathcal{T}_{X,d_D} over $n > 2$ inputs.*

Thm. 12 applies to many loss functions of interest. For example: the average loss function ℓ_{Avg} is strictly monotone in \mathbf{d}_2 although it is non-monotone in \mathbf{d}_D ; ℓ_{bin} is strictly monotone in \mathbf{d}_D but not strictly monotone in \mathbf{d}_2 .

Curiously, Thm. 12 this depends on strictly monotone loss functions wrt *any* metric \mathbf{d} . To give some intuition about this result, the proof (detailed in Appendix §C) shows that strict monotonicity (wrt *any* \mathbf{d}) implies that $\ell \downarrow \{x_1, x_2\}$ is non-trivial for every $x_1 \neq x_2$, and that Kernel mechanisms in the \mathbf{d}_D -privacy space always contain a pair $\{x_1, x_2\}$ for which optimality only holds on trivial loss functions; viz. the result rests on the behaviour of $\mathcal{T}_{X, \mathbf{d}_D}$ mechanisms on 2 secrets.

This suggests that consumers who wish to differentiate over all pairs of secrets will not be well-served by mechanisms of the $\mathcal{T}_{X, \mathbf{d}_D}$ type. However, consumers who wish to distinguish between a *particular* pair of secrets can expect better utility, as optimal mechanisms exist in this case. (This follows from Thm. 10). This is really the nature of the discrete metric – it can distinguish between individual pairs of inputs, but not all pairs at the same time. This highlights the importance of understanding the metrics involved when designing mechanisms with the utility of consumers in mind.

Finally, the next corollary follows from the proof of Thm. 12 (detailed in Appendix §C):

Corollary 13. *There are no universally ℓ_{bin} -optimal mechanisms of type $\mathcal{T}_{X, \mathbf{d}_D}$ for $|X| > 2$.*

This is a strengthening of the result in [5] which proves the above for ε over some threshold. Our result holds for all ε .

V. UNIVERSAL PRIVACY CAPACITIES

§IV-E teaches us that universal ℓ -optimality can be difficult to obtain for arbitrary privacy types. In spite of this it would still be helpful to have some understanding about the ability of the mechanism to transmit useful information whilst still satisfying its privacy constraints. We introduce a weaker universal property of called *privacy type capacity* which we show is always well-defined when there are only finitely-many secrets. The privacy type capacity is an upper bound on the leakage of any mechanism in the type, and therefore provides a universal benchmark for maximal accuracy.

Definition 8. (*Privacy type capacity*) *Given a privacy type $\mathcal{T}_{X, \mathbf{d}}$ we define multiplicative and additive privacy type capacities as follows:*

$$\mathcal{ML}^\times(\mathcal{T}_{X, \mathbf{d}}) := \sup_{C \in \mathcal{T}_{X, \mathbf{d}}} (\mathcal{ML}^\times(C)) \quad (16)$$

and

$$\mathcal{ML}^+(\mathcal{T}_{X, \mathbf{d}}) := \sup_{C \in \mathcal{T}_{X, \mathbf{d}}} (\mathcal{ML}^+(C)) . \quad (17)$$

Theorem 14 (Privacy type capacity is well defined). *Given any privacy type $\mathcal{T}_{X, \mathbf{d}}$, if $|X|$ is finite then both multiplicative and additive type capacities exist (Def. 8) and for each, there exists mechanisms in $\mathcal{T}_{X, \mathbf{d}}$ which achieve the capacity bound.*

Proof: From the fundamental characterisation Thm. 3 every mechanism in $\mathcal{T}_{X, \mathbf{d}}$ is a refinement of a convex sum of

Kernel mechanisms. When $|X|$ is finite there are only finitely many of these. This means that the maximal capacity amongst Kernel mechanisms dominates the capacity of an convex sum of them, or refinement of them. The result now follows. ■

The next result shows the relationship between type capacity and universal optimality results: since ℓ_{bin} and ℓ_{nib} are the gain functions that compute the capacity, if they belong to a mechanism's associated set of universal optimality gain functions then they will achieve the privacy type bound.

Theorem 15 (Universal optimal capacity). *Given any privacy type $\mathcal{T}_{X, \mathbf{d}}$, if $\ell_{\text{bin}} \in L_M^\mathcal{T}$ then $\mathcal{ML}^\times(M) \geq \mathcal{ML}^\times(M')$ for any mechanism $M' \in \mathcal{T}_{X, \mathbf{d}}$. Similarly if $\ell_{\text{nib}} \in L_M^\mathcal{T}$ then $\mathcal{ML}^+(M) \geq \mathcal{ML}^+(M')$ for any mechanism $M' \in \mathcal{T}_{X, \mathbf{d}}$.*

The proof appeals to the additive and multiplicative miracle theorems [3] (restated at Thm. 4 and Thm. 5). Observe that these results give quantitative measurements of the accuracy of mechanisms within a whole privacy type.

In general, though, even when a privacy type does not contain *any* mechanism whose associated optimality loss function classes contain either ℓ_{bin} or ℓ_{nib} , we are still able to compute the capacity and discover a mechanism that achieves the capacity bounds using convex optimisation.

Lemma 16. *The mechanism that achieves the optimal capacity within a privacy type $\mathcal{T}_{X, \mathbf{d}}$ can be written as a channel with no more than $|X|$ columns.*

Corollary 17. *The additive capacity for privacy type $\mathcal{T}_{X, \mathbf{d}}$ where $X = \{x_1, \dots, x_n\}$ is $1 - c$ where c is the optimal value for the linear optimisation problem described by:*

$$\begin{aligned} & \text{minimise } m_{11} + \dots + m_{nn} \\ & \sum_j m_{ij} = 1 \quad 1 \leq i \leq n \\ & m_{ij} \leq m_{kj} \times e^{d(x_i, x_k)} \quad 1 \leq i, k \leq n \\ & m_{ij} \geq 0 \quad 0 \leq i, j \leq n \end{aligned}$$

Similarly we can compute the multiplicative capacity within a type.

Corollary 18. *The multiplicative capacity for privacy type $\mathcal{T}_{X, \mathbf{d}}$ where $X = \{x_1, \dots, x_n\}$ can be found as a linear optimisation problem described by:*

$$\begin{aligned} & \text{maximise } m_{11} + \dots + m_{nn} \\ & \sum_j m_{ij} = 1 \quad 1 \leq i \leq n \\ & m_{ij} \leq m_{kj} \times e^{d(x_i, x_k)} \quad 1 \leq i, k \leq n \\ & m_{ij} \geq 0 \quad 0 \leq i, j \leq n \end{aligned}$$

Observe that for both capacities the optimisation problem contains $O(|X|^2)$ constraints. In some situations this number can be reduced, depending on the properties of the metric.

Note that Cor. 17 and Cor. 18 are both examples of the optimisation formulations similar to that used by Ghosh et al. [4] for loss functions related to “counting queries”. Significant here though is that our results produce universal measurements for privacy type capacity as well as a simplification of how to compute the average minimal loss, yielding the simple linear

objective function given here (rather than minimising over a convex function as in Eqn (1) or (5) which accounts for the “remapping” feature).

VI. EXAMPLES OF PRIVACY TYPES AND THEIR OPTIMALITY RESULTS

In this section we focus on characterisations for some metric spaces of interest, including the two spaces addressed in the literature: $(\mathbb{N}, \mathbf{d}_2)$ and $(\mathcal{X}, \mathbf{d}_D)$ for which the results of Ghosh et al. and Brenner and Nissim hold respectively. We enumerate the kernel mechanisms in each of these spaces by computing the extreme points of the convex region of hypers, and then finding sets of (up to) n posteriors which average to the uniform distribution.

A. Privacy type defined by $(\mathbb{N}, \mathbf{d}_2)$

This privacy type is the one studied by Ghosh et al., corresponding to “counting queries”. Although rich with mechanisms, the Geometric mechanism is the only well-known of this type. It has two instances – the (infinite) Geometric mechanism with outputs over the whole of \mathbb{Z} and the *truncated* Geometric mechanism, in which the output domain matches the input domain (typically a subsequence of \mathbb{N}). These instances in fact have the same leakage properties, ie., they are *equivalent* as channels and thus produce the same hyper-distribution [24].

Definition 9 (Geometric mechanism). *The α -geometric mechanism $G: \mathcal{X} \rightarrow \mathbb{D}\mathbb{Z}$ has the following channel matrix:*

$$G_{x,y} = \frac{1 - \alpha}{1 + \alpha} \cdot \alpha^{\mathbf{d}_2(x,y)}$$

where $\alpha \in (0, 1]$. This mechanism satisfies ε - \mathbf{d}_2 -privacy where \mathbf{d}_2 is the Euclidean metric and $\varepsilon = -\ln \alpha$.

Unsurprisingly (given its optimality properties), we find that the Geometric mechanism is a kernel mechanism.

Lemma 19. *The (infinite/truncated) Geometric mechanism is a kernel mechanism of privacy type $\mathcal{T}_{\mathbb{N}, \mathbf{d}_2}$.*


Figure 4b depicts the space of $\mathcal{T}_{\mathbb{N}, \mathbf{d}_2}$ mechanisms on 3 secrets. On n inputs, this space is constructed from $2 \times (n - 1)$ linear DP constraints resulting in 2^{n-1} vertices – extreme points of the feasible region of posteriors of this type.

Table 4a lists all kernel mechanisms for n up to 6.

Ghosh et al. showed that the Geometric mechanism is universally optimal for the class of monotone loss functions, meaning that it provides the most accuracy (within its type) to Bayesian consumers described by monotone loss functions. However this result does not provide any way to determine exactly *how* accurate the data release can be, even for this optimal mechanism. A QIF analysis is able to give robust measurements for the amount of leakage as expressed by gain or loss functions; combined with Thm. 15 we are now able to give robust measurements for the whole privacy class. In particular since $\ell_{\text{bin}} \in L_G$, we deduce that the multiplicative channel capacity of the Geometric dominates the multiplicative

Dims	Vertices	Kernel Mechanisms	Mult. Capacity	Add. Capacity
2	2	1	1.33	0.33
3	4	2	1.67	0.5
4	8	11	2	0.67
5	16	187	2.33	0.75
6	32	15346	2.67	0.83

(a) Kernel mechanisms and capacities in the privacy type induced by $\ln 2 \cdot \mathbf{d}_2$.



(b) The space of \mathbf{d}_2 -private hypers on 3 inputs $\{0, 1, 2\}$. The Geometric mechanism for $\varepsilon = \ln 2$ consists of the 3 orange vertices (B, C, D).

Fig. 4: Kernel mechanisms in the space of \mathbf{d}_2 -private hypers.

channel capacity for any mechanism in the whole privacy type. Going further we can compute the exact dominating channel capacity. (Refer to Table 4a for numerical examples.)

Corollary 20. (Dominating multiplicative capacity) *Let N be the size of \mathcal{X} and let $\alpha = e^{-\varepsilon}$, where ε is given by the definition of DP. Then for any mechanism $M \in \mathcal{T}_{\mathcal{X}, \mathbf{d}_2}$ we have that its multiplicative capacity is no more than $(N^{(1-\alpha)+2\alpha})/(1+\alpha)$.*

We also observe that the Geometric mechanism is the kernel mechanism which obtains the *additive* capacity for the mechanisms we computed (cf. Table 4a), suggesting that it may also be optimal for the non-monotone ℓ_{nb} loss function.

B. Privacy type defined by $(\mathcal{X}, \mathbf{d}_D)$

This privacy type is the one studied by Brenner and Nissim, corresponding to “sum queries”. For an input space \mathcal{X} of size n , the vertices of the convex region are points of intersection of $n - 1$ hyperplanes, corresponding to $n - 1$ distinct privacy constraints. Since there are only 2 possible distances under the discrete metric, each vertex can only contain 2 possible values, and so the number of possible vertices is given by $\sum_{i=1}^{n-1} \binom{n}{i} = 2^n - 2$.⁵ We present some of these mechanisms in Table 5a, including the well-known “Random Response” mechanism.

Definition 10 (Random Response mechanism). *The α -random response mechanism $R: \mathcal{X} \rightarrow \mathbb{D}\mathcal{X}$ has the following channel*

⁵Thinking of each vertex as having either the ‘high’ or the ‘low’ value, this is the number of combinations for each choice of ‘high’ values.


matrix:

$$\begin{aligned} R_{x,x} &= 1/k \\ R_{x,y} &= \alpha/k \quad \text{for } x \neq y \end{aligned}$$

where k is a normalisation term and $\alpha \in (0, 1]$. This mechanism satisfies ε - \mathbf{d}_D -privacy where \mathbf{d}_D is the discrete metric and $\varepsilon = -\ln \alpha$.

Dims	Vertices	Kernel Mechanisms	Mult. Capacity	Add. Capacity
2	2	1	1.33	0.33
3	6	5	1.5	0.4
4	14	41	1.6	0.43
5	30	1291	1.67	0.44

(a) Kernel mechanisms and capacities in the privacy type induced by $\ln 2 \cdot \mathbf{d}_D$.



(b) The space of \mathbf{d}_D -private hypervolumes over 3 inputs. The Random Response mechanism corresponds to the 3 orange vertices (A, B, C). Its dual R° corresponds to the blue vertices.

Fig. 5: Kernel mechanisms in the space of \mathbf{d}_D -private hypervolumes.

We note that the Random Response mechanism is a kernel mechanism. The space of mechanisms on 3 secrets for the Discrete metric is depicted in Figure 5b.

Lemma 21. *The Random Response mechanism is a kernel mechanism in the privacy type $\mathcal{T}_{\mathcal{X}, \mathbf{d}_D}$.*

Brenner and Nissim studied mechanisms of this privacy type, focussing on Bayesian consumers and ℓ_{bin} , showing that no mechanism is universally optimal for ℓ_{bin} . Their result demonstrates that there is no mechanism for the Discrete privacy type with the same universal optimality properties as the Geometric mechanism in the Euclidean privacy type. However Thm. 10 indicates that there are in fact mechanisms which are ℓ -optimal for some non-trivial (monotone) loss functions. The next example of a mechanism (also described by Brenner and Nissim) is in fact ℓ -optimal for a non-trivial monotone loss function. Consider the following hypervolumes of type $\mathcal{T}_{\mathcal{X}, \ln 2 \cdot \mathbf{d}_D}$ on n inputs:

$$\Delta_{K_A} = \begin{bmatrix} 1/4 & 2/5 \\ \vdots & \vdots \\ 1/4 & 2/5 \\ 1/2 & 1/5 \\ 4/9 & 5/9 \end{bmatrix} \quad \Delta_{K_B} = \begin{bmatrix} 1/2 & 1/5 \\ 1/4 & 2/5 \\ \vdots & \vdots \\ 1/4 & 2/5 \\ 4/9 & 5/9 \end{bmatrix}$$

We use Thm. 6(5) as follows. Concentrating on inputs $X^\dagger := \{x_1, x_n\}$, we see that $K_A \downarrow X^\dagger = K_B \downarrow X^\dagger = T$ (where T is the universally optimal mechanism from Cor. 9). Since T is the unique Kernel mechanism in a 2-element state space, we have that both are ℓ_e -optimal for loss function $\ell_e(w, x)$ which has two actions, labelled by X^\dagger , and is 0 whenever $w \neq x$, and 1 if $w = x$.⁶ From Figure 3 we can see immediately that ℓ_e is non-trivial, thus by Thm. 6(5) we see that both K_A, K_B are $\ell_e \uparrow \mathcal{X}$ -optimal, as well as all probabilistic combinations $K_A \oplus_p K_B$ defined above.

While we can compute a capacity for the whole Euclidean type from the universal optimality of the Geometric, this is not possible on the Discrete space owing to the impossibility results for ℓ_{bin} . We can use instead the characterisation of its structure directly to compute dominating additive and multiplicative capacities giving a tight quantitative upper bound for the accuracy of any mechanisms in the type. The dominating multiplicative capacity is given by the random response R , and the dominating additive capacity is given by its dual R° as follows. Let $\varepsilon = -\ln \alpha = \ln \beta$, and $k = 1 + (|\mathcal{X}| - 1)e^{-\varepsilon}$ and $m = 1 + (|\mathcal{X}| - 1)e^\varepsilon$, then:

$$\begin{aligned} R_{ij} &= 1/k, \text{ if } i = j \\ &= \alpha/k, \text{ otherwise} \end{aligned} \quad \text{and} \quad \begin{aligned} R^\circ_{ij} &= 1/m, \text{ if } i = j \\ &= \beta/m, \text{ otherwise} \end{aligned} .$$

The dual R° has the minimum elements on the diagonal whereas the Random Response R has its maxima there. (Recall Figure 5b for illustration.)

Corollary 22 (Dominating capacities). *For any M in the Discrete privacy type, we have $\mathcal{ML}^\times(M) \leq |\mathcal{X}| / (1 + (|\mathcal{X}| - 1)e^{-\varepsilon})$ and $\mathcal{ML}^+(M) \leq 1 - 1 / (1 + (|\mathcal{X}| - 1)e^\varepsilon)$.*

Proof. From Thm. 3, M must be a refinement of a convex sum of vertex mechanisms. From the set of vertices we can see that any inner δ in the support of $[v \triangleright M]$ must satisfy $1/m \leq \delta_x \leq 1/k$. The result now follows from the definition of R and R° and the miracle theorems Thm. 4 and Thm. 5.

C. Privacy type defined by $(\mathcal{X}, \mathbf{d}_2)$ on grids

This privacy type is described by the Euclidean distance \mathbf{d}_2 on $n \times n$ grids, as might be used in geo-location privacy, for example. Here we find that there are a large number of vertices on even the 2×2 space, making computing all of the kernel mechanisms expensive. Nevertheless, using Cor. 17 and Cor. 18 we can efficiently compute the capacities even when the kernel mechanisms cannot be enumerated. Table 6a illustrates some results for 3 grids.

In addition, optimality results in this space can be obtained for computed kernel mechanisms using Thm. 6(5) by employing the same technique illustrated in §VI-B.


D. Privacy type on the Hamming cube

This privacy type is defined by the Hamming distance \mathbf{d}_H on bitstrings. As with other examples, the dimensions of this type is the number of rows in the mechanism, ie., 2^n for bitstrings of

⁶In fact both $K_A \downarrow X^\dagger, K_B \downarrow X^\dagger$ are ℓ -optimal for all loss functions on X^\dagger

Dims	Vertices	Kernel Mechanisms	Mult. Capacity	Add. Capacity
1×1	18	403	1.68	0.48
2×2	4798	> 10000	2.5	0.62
3×3	-	-	3.53	0.79

(a) Kernel mechanisms in the $\mathcal{T}_{\mathcal{X},d_2}$ privacy type on grids.




(b) The $\mathcal{T}_{\mathcal{X},d_2}$ type on a 2x2 grid.

Fig. 6: The Euclidean privacy type on grids.

length n , and the vertices are the extreme points of the feasible region of posteriors. Here, even on bitstrings of length 3 the number of kernel mechanisms exceeds 29000, although we again are able to compute capacities for higher dimensions. These are illustrated in Table 7a.

Bits	Dims	Vertices	Kernel Mechanisms	Mult. Capacity	Add. Capacity
2	4	6	4	1.78	0.56
3	8	38	29275	2.37	0.70
4	16	-	-	3.16	0.80

(a) Kernel mechanisms in the $\mathcal{T}_{\mathcal{X},d_H}$ privacy type for $\mathcal{X} = \{0, 1\}^3$.



(b) The Hamming type on bitstrings of length 3. The orange lines depict a path on which the d_H metric is linear.

Fig. 7: The privacy type defined by the Hamming cube.

Figure 8 depicts a kernel mechanism K and its associated hyper Δ_K for the Hamming cube on 3 secrets. We now show how to construct a loss function ℓ for which K is universally ℓ -optimal.

Observe (Figure 8) that for $X = \{000, 100, 110, 111\}$, the mechanism $K \downarrow X$ is exactly the Geometric mechanism G (for $\varepsilon = \ln 2$). In other words, ℓ_{bin} is in the utility set $L_G^{\mathcal{T}_{\mathcal{X},d_2}}$ which coincides with $L_G^{\mathcal{T}_{\mathcal{X},d_H}}$ (noting that the Hamming metric is linear on the secrets X). Therefore, from Thm. 6(5) we have that $\ell_{\text{bin}} \uparrow \mathcal{X}$ is in $L_K^{\mathcal{T}_{\mathcal{X},d_H}}$, using the fact that $K \downarrow X = G$. Therefore K is universally $\ell_{\text{bin}} \uparrow \mathcal{X}$ -optimal. And in fact K will be universally $\ell \uparrow \mathcal{X}$ -optimal for any monotone loss function ℓ defined on X . Such a loss function would be useful to a

consumer who is only interested in learning the inputs X and finds no value in the remaining inputs $\mathcal{X} \setminus X$.

Note that $X = \{000, 100, 110, 111\}$ corresponds to the orange path depicted in Figure 7b for which we observed that the Hamming metric is linear.

$\Delta_K =$	K	y_1	y_2	y_3	y_4
		000	2/3	1/6	1/12
100	1/3	1/3	1/6	1/6	
110	1/6	1/6	1/3	1/3	
111	1/12	1/12	1/6	2/3	
010	1/3	1/3	1/6	1/6	
011	1/6	1/6	1/3	1/3	
101	1/6	1/6	1/3	1/3	
001	1/3	1/3	1/6	1/6	

Fig. 8: A kernel hyper (Δ_K) and corresponding kernel mechanism (K) on bitstrings of length 3 wrt the Hamming distance d_H with $\varepsilon = \ln 2$. K is universally ℓ -optimal for many loss functions derived from the optimality of the Geometric mechanism on 4 secrets.

VII. CONCLUSION

We have extended the notion of universal optimality to d -privacy. Our principal technique is an interpretation of QIF, enabling a geometric interpretation of a class of privacy mechanisms and a characterisation of privacy types in terms of refinement. A principal contribution is to clarify the negative results of Brenner and Nissim, extending them beyond ℓ_{bin} to other strictly monotonic mechanisms, and removing the connection to ε . Our study of the relationship of utility sets provides techniques for finding new optimality results to new domains beyond sum and counting queries, as illustrated by our example of mechanisms over the hamming cube. Finally where the underlying domain is very complex, making it unlikely to find universally optimal results we introduce the notion of privacy type capacity and show how to compute it for both multiplicative and additive capacities. With these results we provide strong benchmarks for leakage measurements in domains where full formal analysis is difficult.

REFERENCES

- [1] C. Dwork, F. Mcsherry, K. Nissim, and A. Smith, “Calibrating noise to sensitivity in private data analysis,” in *In Proceedings of the Third Theory of Cryptography Conference (TCC)*, ser. Lecture Notes in Computer Science, S. Halevi and T. Rabin, Eds., vol. 3876. Springer, 2006, pp. 265–284.
- [2] C. Dwork, “Differential privacy,” in *33rd International Colloquium on Automata, Languages and Programming (ICALP 2006)*, ser. Lecture Notes in Computer Science, M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, Eds., vol. 4052. Springer, 2006, pp. 1–12. [Online]. Available: http://dx.doi.org/10.1007/11787006_1
- [3] M. S. Alvim, K. Chatzikokolakis, A. McIver, C. Morgan, C. Palamidessi, and G. Smith, “The science of quantitative information flow,” 2019.
- [4] A. Ghosh, T. Roughgarden, and M. Sundararajan, “Universally utility-maximizing privacy mechanisms,” *SIAM Journal on Computing*, vol. 41, no. 6, pp. 1673–1693, 2012.

[5] H. Brenner and K. Nissim, “Impossibility of differentially private universally optimal mechanisms,” in *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. IEEE, 2010, pp. 71–80.

[6] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, “Local privacy and statistical minimax rates,” in *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE Computer Society, 2013, pp. 429–438. [Online]. Available: <https://doi.org/10.1109/FOCS.2013.53>

[7] K. Chatzikokolakis, M. E. Andrés, N. E. Bordenabe, and C. Palamidessi, “Broadening the scope of differential privacy using metrics,” in *International Symposium on Privacy Enhancing Technologies Symposium*. Springer, 2013, pp. 82–102.

[8] M. E. Andrés, N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi, “Geo-indistinguishability: Differential privacy for location-based systems,” in *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, 2013, pp. 901–914.

[9] N. Fernandes, M. Dras, and A. McIver, “Generalised differential privacy for text document processing,” in *Principles of Security and Trust - 8th International Conference, POST 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings*, 2019, pp. 123–148. [Online]. Available: https://doi.org/10.1007/978-3-030-17138-4_6

[10] N. Fernandes, A. McIver, and C. Morgan, “The laplace mechanism has optimal utility for differential privacy over continuous queries,” in *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*. IEEE, 2021, pp. 1–12. [Online]. Available: <https://doi.org/10.1109/LICS52264.2021.9470718>

[11] M. Gupte and M. Sundararajan, “Universally optimal privacy mechanisms for minimax agents,” in *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, 2010, pp. 135–146.

[12] J. Acharya, K. Bonawitz, P. Kairouz, D. Ramage, and Z. Sun, “Context aware local differential privacy,” in *International Conference on Machine Learning*. PMLR, 2020, pp. 52–62.

[13] P. Kairouz, S. Oh, and P. Viswanath, “Extremal mechanisms for local differential privacy,” *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 492–542, 2016.

[14] F. Koufogiannis, S. Han, and G. J. Pappas, “Optimality of the laplace mechanism in differential privacy,” *arXiv preprint arXiv:1504.00065*, 2015.

[15] H. A. J. C. Duchi, “Near instance-optimality in differential privacy,” 2020, arXiv:2005.10630v1, 2020.

[16] M. S. Alvim, K. Chatzikokolakis, C. Palamidessi, and G. Smith, “Measuring information leakage using generalized gain functions,” in *2012 IEEE 25th Computer Security Foundations Symposium*. IEEE, 2012, pp. 265–279.

[17] A. McIver, C. Morgan, G. Smith, B. Espinoza, and L. Meinicke, “Abstract channels and their robust information-leakage ordering,” in *Principles of Security and Trust - Third International Conference, POST 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings*, ser. Lecture Notes in Computer Science, M. Abadi and S. Kremer, Eds., vol. 8414. Springer, 2014, pp. 83–102.

[18] S. T. Rachev and L. Rüschendorf, *Mass Transportation Problems: Volume I: Theory*. Springer Science & Business Media, 1998, vol. 1.

[19] K. Chatzikokolakis, N. Fernandes, and C. Palamidessi, “Comparing systems: Max-case refinement orders and application to differential privacy,” in *Proc. CSF*. IEEE Press, 2019.

[20] M. S. Alvim, M. E. Andrés, K. Chatzikokolakis, and C. Palamidessi, “On the relation between differential privacy and quantitative information flow,” in *Automata, Languages and Programming - 38th International Colloquium, ICALP 2011, Zurich, Switzerland, July 4-8, 2011, Proceedings, Part II*, 2011, pp. 60–76. [Online]. Available: https://doi.org/10.1007/978-3-642-22012-8_4

[21] N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi, “Optimal geo-indistinguishable mechanisms for location privacy,” in *Proc. CCS*, 2014, pp. 251–262.

[22] M. S. Alvim, N. Fernandes, A. McIver, and G. H. Nunes, “On privacy and accuracy in data releases (invited paper),” in *31st International Conference on Concurrency Theory, CONCUR 2020, September 1-4, 2020, Vienna, Austria (Virtual Conference)*, ser. LIPIcs, I. Konnov and L. Kovács, Eds., vol. 171. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020, pp. 1:1–1:18. [Online]. Available: <https://doi.org/10.4230/LIPIcs.CONCUR.2020.1>

[23] Úlfar Erlingsson, V. Pihur, and A. Korolova, “RAPPOR: Randomized aggregatable privacy-preserving ordinal response,” in *Proc. CCS*, 2014, pp. 1054–1067.

[24] K. Chatzikokolakis, N. Fernandes, and C. Palamidessi, “Refinement orders for quantitative information flow and differential privacy,” *Journal of Cybersecurity and Privacy*, vol. 1, no. 1, pp. 40–77, 2021.

APPENDIX

A. Results supporting §II

The following shows the equivalence between the formulation of utility loss in the literature ([4], [5]) and the QIF formulation of posterior uncertainty.

Lemma A1. *The expected utility loss in Eqn 1 is equivalent to the posterior ℓ -uncertainty defined in Eqn 5.*

Proof: We reason as follows:

$$\begin{aligned}
 & U_\ell(\pi, M) \\
 = & \min_r \sum_{x \in \mathcal{X}} \pi_x \sum_{y \in \mathcal{Y}} M_{x,y} \ell(r(y), x) && \text{“Eqn 1”} \\
 = & \min_r \sum_{y \in \mathcal{Y}} \sum_{x \in \mathcal{X}} \pi_x M_{x,y} \ell(r(y), x) && \text{“Reorganising”} \\
 = & \sum_{y \in \mathcal{Y}} \min_r \sum_{x \in \mathcal{X}} \pi_x M_{x,y} \ell(r(y), x) && \text{“Min } r \text{ is over each } y\text{”} \\
 = & \sum_{y \in \mathcal{Y}} \min_w \sum_{x \in \mathcal{X}} \pi_x M_{x,y} \ell(w, x) && \text{“Letting } w = r(y)\text{”} \\
 = & U_\ell[\pi \triangleright M] && \text{“Eqn 5”}
 \end{aligned}$$

■

Note that in the second-last step we employ a mapping from remaps $r(y)$ to guesses w . We can think of this as simply a relabelling of observations y to guesses w , and we can always ‘duplicate’ guesses w so that there is a one-one correspondence between guesses and inputs.

B. Results supporting §III

The following establishes the correspondence between \mathbf{d} -private mechanisms C and hyper-distributions.

Lemma 2. *C is \mathbf{d} -private iff $\delta_x^y - e^{\mathbf{d}(x,x')} \delta_{x'}^y \leq 0$ for every δ^y in the support of $[v \triangleright C]$ where $v \in \mathbb{D}\mathcal{X}$ is the uniform distribution.*

Proof: Recall that C is \mathbf{d} -private whenever $\frac{C_{x,y}}{C_{x',y}} \leq e^{\mathbf{d}(x,x')}$. Now rewrite δ_x^y as $v_x \cdot C_{x,y} / \sum_{z \in \mathcal{X}} v_x \cdot C_{z,y}$ and the result follows. ■

We next set out some key properties of kernel mechanisms which allow us to establish the fundamental characterisation. The first property tells us that kernel mechanisms are *irreducible*.

Property 1. *If K is a kernel mechanism then there is no kernel mechanism K' st. $[\Delta_{K'}] \subset [\Delta_K]$ where $\Delta_K, \Delta_{K'}$ are hypers corresponding to mechanisms K, K' respectively.*

Proof: Assume that K' is a kernel mechanism and $K' \subset K$ (treating K, K' as sets of posteriors). Then the posteriors of K' must be linearly independent (from Def. 5). Therefore the uniform distribution can be uniquely expressed as a linear combination of posteriors of K' . ie. We can write $u = \sum_i a_i \delta^i$ for scalars a_i and posteriors δ^i . But the same is true for K . ie.

Its posteriors are linearly independent and there is a unique combination that produces the uniform distribution. ie. We write $u = \sum_j b_j \gamma^j$. But the posteriors of $K' \subset K$. Therefore we can write

$$\begin{aligned} & \sum_i a_i \delta^i \\ = & \sum_j b_j \gamma^j && \text{“Both equal to } u\text{”} \\ = & \sum_i b_i \delta^i + \sum_{k=j-i} b_k \gamma^k && \text{“} K' \subset K \text{”} \end{aligned}$$

But if any of the scalars b_k is non-zero, then we have 2 unique linear combinations of γ^j which average to the uniform distribution (one involving only a_i and δ^i and the other involving the b_k and γ^k), contradicting the linear independence of the posteriors γ^j . Therefore we must have $b_k = 0$ and $b_i = a_i$ and so $K' = K$, contradicting $K' \subset K$. Thus no such K' can exist. ■

The next property says that kernel mechanisms generate all of the vertex mechanisms.

Property 2. *Any vertex mechanism can be written (non-uniquely) as a convex sum of kernel mechanisms. Conversely, any convex sum of kernel mechanisms is a vertex mechanism.*

Proof: For the converse, note that the convex sum of hypers takes the union of posteriors, which is clearly a vertex mechanism. For the forward direction, we describe an algorithm whose invariant is the expected value E of the remaining posteriors of M (wrt the remaining outers of M). We start by removing probability mass corresponding to any kernel mechanism M_1 whose set of posteriors is a subset of the set of posteriors of M . Since M_1 is a valid mechanism, it has the same expected value E as M , so this operation preserves E . We continue along this vein for every kernel mechanism we can find. Finally if we are left with some set of weighted posteriors, whose expected value must also be E and so they must form a kernel mechanism, utilising the rest of the probability mass. ■

Our final property says that kernel mechanisms are not related by refinement.

Property 3. *If K, K^* are kernel mechanisms then $K \not\sqsubseteq K^*$ and $K^* \not\sqsubseteq K$.*

Proof: Let Δ, Δ^* be the hypers of K, K^* respectively. From Property 1 we cannot have $\Delta \subset \Delta^*$ or $\Delta^* \subset \Delta$ (considering Δ, Δ^* as sets of posteriors). Therefore, by convexity of the region, the posteriors of Δ cannot lie inside the convex hull of posteriors of Δ^* and vice versa. The result follows. ■

Observe that although Property 2 says that kernel mechanisms generate the space of vertex mechanisms, the representation of a vertex hyper as a convex sum of kernel hypers is not necessarily unique.

We are now ready to prove the characterisation of $\mathcal{T}_{X,d}$ mechanisms.

Theorem 3. *(Fundamental Characterisation of $\mathcal{T}_{X,d}$ Mechanisms) Every $\mathcal{T}_{X,d}$ mechanism is a refinement of a convex sum of kernel mechanisms.*

Proof: Given $M \in \mathcal{T}_{X,d}$ which is not a vertex mechanism, and its corresponding hyper Δ_M , we know that each posterior δ^i in $[\Delta_M]$ sits inside the convex hull of the vertices in the space of \mathbf{d} -private hypers, and we can thus perform a “reverse Earth Move” (recall Observation a), moving mass from each posterior δ^i to some set of vertices whose convex hull encloses δ^i , preserving the overall centre of mass ν of Δ . Thus we get a valid anti-refining vertex mechanism, which from Property 2 is a convex sum of kernel mechanisms. ■

C. Results supporting §IV

In this section we detail the proofs of Thm. 6 and Thm. 10 as well as the impossibility result of Thm. 11 which states that there are no universally optimal mechanisms on $n > 2$ inputs. We also tease out the details of the impossibility result for ℓ -optimal mechanisms in Thm. 12.

We begin with the result of Thm. 6 which lays out the duality between loss functions and mechanisms.

Theorem 6 (ℓ -Optimal Duality). *Let $\mathcal{T}_{X,d}$ and $\overline{\mathcal{T}}_{X,d}$ be privacy types where $X \subseteq \mathcal{X}$. For mechanisms $M, M' \in \mathcal{T}_{X,d}$, we have the following.*

- 1) $\mathcal{L}^* \subseteq L_M^{\overline{\mathcal{T}}_{X,d}}$,
- 2) $M \sqsubseteq M'$ implies $L_M^{\overline{\mathcal{T}}_{X,d}} \supseteq L_{M'}^{\overline{\mathcal{T}}_{X,d}}$,
- 3) $L_{\mathbb{I}}^{\overline{\mathcal{T}}_{X,d}} = \mathcal{L}^*$, where \mathbb{I} is the trivial mechanism,
- 4) $L_{M \oplus_p M'}^{\overline{\mathcal{T}}_{X,d}} = L_M^{\overline{\mathcal{T}}_{X,d}} \cap L_{M'}^{\overline{\mathcal{T}}_{X,d}}$ for $0 < p < 1$,
- 5) $\ell \in L_{M \downarrow X}^{\overline{\mathcal{T}}_{X,d}}$ if and only if $\ell \uparrow X \in L_M^{\overline{\mathcal{T}}_{X,d}}$.

Proof:

- 1) $\mathcal{L}^* \subseteq L_M^{\overline{\mathcal{T}}_{X,d}}$. This follows since for $\ell \in \mathcal{L}^*$ we have that $U_\ell(\pi, M) = U_\ell(\pi)$ for any mechanism M . Therefore all mechanisms are ℓ -optimal for trivial loss functions.
- 2) $M \sqsubseteq M'$ implies $L_M^{\overline{\mathcal{T}}_{X,d}} \supseteq L_{M'}^{\overline{\mathcal{T}}_{X,d}}$. Let $\ell \in L_{M'}^{\overline{\mathcal{T}}_{X,d}}$. We have that:

$$U_\ell(\pi, M) \leq U_\ell(\pi, M') \leq U_\ell(\pi, M),$$

where the first inequality holds by the refinement assumption, and the second by the universal ℓ -optimality of M' . Thus $U_\ell(\pi, M) = U_\ell(\pi, M')$ and therefore $\ell \in L_M^{\overline{\mathcal{T}}_{X,d}}$.

- 3) $L_{\mathbb{I}}^{\overline{\mathcal{T}}_{X,d}} = \mathcal{L}^*$, where \mathbb{I} is the trivial mechanism.

By Lem. 1 and Thm. 3 we can represent the hyperdistribution $[\nu, K]$ as the sum $\sum_i a_i v_i$, where v_i are vertices in the convex space. Suppose that ℓ is not trivial and that $U_\ell(\nu) = \sum_x \ell(x, w^*)$, then at least one vertex v must have $U_\ell(v) = \sum_x \ell(x, w)$, where $w \neq w^*$. Assume wlog. that $v = v_1$. This means that:

$$\begin{aligned} & U_\ell(v, \mathbb{I}) \\ = & U_\ell(v) \\ = & \sum_i a_i U_\ell(v) \\ = & \sum_i a_i (\sum_x \ell(x, w^*)) \\ > & \sum_i a_i \min_{w \in \mathcal{W}} (\sum_x \ell(x, w)) \\ = & U_\ell(v, K), \end{aligned}$$

thus $\ell \notin L_{\mathbb{I}}^{\mathcal{T}_{X,d}}$.

4) $L_{M \oplus_p M'}^{\mathcal{T}_{X,d}} = L_M^{\mathcal{T}_{X,d}} \cap L_{M'}^{\mathcal{T}_{X,d}}$ for $0 < p < 1$. Note that $U_\ell(\pi, M \oplus_p M') = p \times U_\ell(\pi, M) + (1-p) \times U_\ell(\pi, M')$. Thus if $\ell \in L_M^{\mathcal{T}_{X,d}} \cap L_{M'}^{\mathcal{T}_{X,d}}$ it must also be contained in $L_{M \oplus_p M'}^{\mathcal{T}_{X,d}}$.

Conversely, if $\ell \in L_M^{\mathcal{T}_{X,d}} - L_{M'}^{\mathcal{T}_{X,d}}$ then there is some π such that $U_\ell(\pi, M) < U_\ell(\pi, M')$, in which case $U_\ell(\pi, M) < U_\ell(\pi, M \oplus_p M')$ also, since $0 < p < 1$.

5) $\ell \in L_{M \downarrow X}^{\mathcal{T}_{X,d}}$ if and only if $\ell \uparrow \mathcal{X} \in L_M^{\mathcal{T}_{X,d}}$. This follows since for $\pi \in \mathbb{D}X$ we have $U_\ell(\pi, M \downarrow X) = U_{\ell \uparrow \mathcal{X}}(\pi, M)$. More generally $U_{\ell \uparrow \mathcal{X}}(\pi, M) = \alpha \times U_\ell(\pi', M \downarrow X)$, where $\alpha = \pi(X)$ and $\pi'_x = \pi_x / \alpha$ for $x \in X$.

Next we detail the proof of universal optimality existence, which relies on the key idea that optimality on a subset of inputs can be extended to optimality on the full mechanism.

Theorem 10 (Universal optimality existence). *Given any privacy type $\mathcal{T}_{X,d}$, if $|X| > 1$ then there is a non-trivial monotone loss function ℓ and a mechanism $M \in \mathcal{T}_{X,d}$ such that M is universally ℓ -optimal.*

Proof: Our proof is constructive - we show how to construct an ℓ -optimal mechanism for any \mathbf{d} using a carefully crafted ℓ .

For any discrete metric space $(\mathcal{X}, \mathbf{d})$, choose $x_1^*, x_2^* \in \mathcal{X}$ such that $\mathbf{d}(x_1^*, x_2^*) \leq \mathbf{d}(x_1, x_2)$ for all $x_1, x_2 \in \mathcal{X}$. Then construct the following hyper:

Δ	$k/k+j$	$j/k+j$
x_1^*	$1/k$	α/j
x_2^*	α/k	$1/j$
x_a	α/k	$1/j$
x_b	α/k	$1/j$
\dots	\dots	\dots
x_n	α/k	$1/j$

where $\alpha = e^{-\epsilon}$ and k, j are normalising constants ensuring columns sum to 1.

It is easy to check that the posteriors average to the uniform distribution using the outer probabilities, hence the hyper corresponds to a proper mechanism K . We now observe that K is \mathbf{d} -private for any \mathbf{d} . For any $x_1, x_2 \neq x_1^*, x_2^*$ the \mathbf{d} -privacy constraints are trivially satisfied in Δ and therefore in K (by Lem. 2). Also, $\delta_{x_1^*}^y \leq \delta_{x_2^*}^y \times e^{\mathbf{d}(x_1^*, x_2^*)} \leq \delta_{x_2^*}^y \times e^{\mathbf{d}(x_1^*, x_2^*)}$ (by assumption) and similarly for x_2^* . Therefore the \mathbf{d} -privacy constraints are satisfied on Δ and so by Lem. 2, K is \mathbf{d} -private.

Now, if we construct $K \downarrow \{x_1^*, x_2^*\}$ we get:

$K \downarrow \{x_1^*, x_2^*\}$	y_1	y_2
x_1^*	$1/(k+j)$	$\alpha/(k+j)$
x_2^*	$\alpha/(k+j)$	$1/(k+j)$

which corresponds to the universally optimal mechanism on 2 inputs (Cor. 9). Choosing any ℓ on 2 inputs, we can lift it to a

loss function $\ell \uparrow \mathcal{X}$ on the whole domain, which by Thm. 6(5) is therefore universally optimal for K . ■

Finally we detail the proofs of the impossibility results. In this section we will denote by T the universally optimal mechanism on 2 secrets (cf. Cor. 9). We will say that $K \equiv T$ whenever K and T represent the same abstract channel.

We first prove that for the space of $n > 2$ inputs there is no single minimal element under refinement.

Lemma A2. *Let (X, d) be a metric space and let $|X| > 2$. Then the space of d -hypers contains at least $2(n-1)$ vertices.*

Proof: Pick any 3 inputs x_1, x_2, x_3 . If these inputs are collinear (ie. $d(x_1, x_2) + d(x_2, x_3) = d(x_1, x_3)$) then the \mathbf{d} -privacy constraints on (x_1, x_2) and (x_2, x_3) imply the constraints on (x_1, x_3) , and so the 3 hyperplanes (corresponding to constraints in each direction) contribute only 4 vertices to the convex region. If the points are not collinear, then each pair (x_i, x_j) contributes 2 constraints, resulting in 6 constraint hyperplanes in total and 6 points in the convex space. The minimum number of vertices in the space therefore occurs when the set X is totally ordered (ie. maximising the number of linear, and thus inferred, constraints), and so we need only consider constraints on ‘adjacent’ vertices. In total this yields $2(n-1)$ vertices (corresponding to 2 constraints per adjacent pair). ■

We now have the details in place to prove the impossibility of universally optimal mechanisms for $n > 2$.

Theorem 11 (Impossibility of Universally Optimal Mechanisms). *There are no universally optimal mechanisms in $\mathcal{T}_{X,d}$ for any \mathbf{d} and for $|X| > 2$.*

Proof: By Lem. A2 we have that the space contains at least $2(n-1)$ vertices. Since this is larger than n for $n > 2$, we must have more than one kernel mechanism, and thus more than one vertex mechanism. The result follows. ■

To prove Thm. 12 we need a series of technical lemmas. We will first show that strictly monotone ℓ implies its restriction to pairs of secrets is non-trivial (Lem. A3). We then show that the latter condition implies the impossibility (Lem. A6).

We begin by observing that trivial loss functions can be defined as follows:

Definition A1 (Trivial Loss Function). *A loss function $\ell : \mathcal{W} \times \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ is called trivial if there exists a $w^* \in \mathcal{W}$ such that $\ell(w^*, x) \leq \ell(w, x)$ for all $x \in \mathcal{X}$ and $w \in \mathcal{W} \setminus \{w^*\}$.*

We now have the following property of strictly monotone loss functions.

Lemma A3. *Given a loss function $\ell : \mathcal{W} \times \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$, if there exists a metric \mathbf{d} such that ℓ is strictly monotone in \mathbf{d} , then $\ell \downarrow \{x, x'\}$ must be non-trivial for all $x, x' \in \mathcal{X}$ with $x \neq x'$.*

Proof: Assume that there exists $x, x' \in \mathcal{X}$ with $x \neq x'$ st. $\ell \downarrow \{x, x'\}$ is trivial. Then by Def. A1 there exists a w^* st. $\ell(w^*, x) \leq \ell(w, x)$ and $\ell(w^*, x') \leq \ell(w, x')$ for all $w \neq w^*$. Since ℓ is strictly monotone we must have $\mathbf{d}(\alpha(w^*), x) \leq$

$\mathbf{d}(\alpha(w), x)$ and $\mathbf{d}(\alpha(w^*), x') \leq \mathbf{d}(\alpha(w), x')$ for all $w \neq w^*$. But we also know from strict monotonicity that there exists $w_a, w_b \in \mathcal{W}$ with $w_a \neq w_b$ st. $\alpha(w_a) = x$ and $\alpha(w_b) = x'$. Therefore we have

$$\begin{aligned} \mathbf{d}(\alpha(w^*), x) &\leq \mathbf{d}(\alpha(w_a), x) = 0 \\ \mathbf{d}(\alpha(w^*), x') &\leq \mathbf{d}(\alpha(w_b), x') = 0 \end{aligned}$$

Since α is injective, this implies that $w_a = w_b$ and hence $x = x'$, contradicting our assumption that $x \neq x'$. Therefore no such x, x' exists. ■

Next we need the following existence theorem and its corollary which applies to the \mathbf{d}_D space.

Theorem A4. *In the space of \mathbf{d}_D -private hyps there exist non-trivial monotone (on \mathbf{d}_D) loss functions \mathcal{L} and mechanisms which are universally \mathcal{L} -optimal.*

Proof: The proof uses the same construction as for Thm. 10, which we note must be a \mathbf{d}_D -private mechanism. We choose the following loss function, a lifting of an optimal loss function on 2 inputs:

ℓ	w_1	w_2
x_1	0	1
x_2	1	0
x_3	0	0
x_4	0	0
...
x_n	0	0

Since $\ell \downarrow \{x_1, x_2\}$ is monotone on \mathbf{d}_D , we have that ℓ is monotone on \mathbf{d}_D and the mechanism K (with corresponding hyper Δ) is universally ℓ -optimal. ■

The following corollary explains why we can always find universally optimal mechanisms in this space.

Corollary A5. *Given any pair of inputs $x_1, x_2 \in \mathcal{X}$, there are always \mathbf{d}_D -private mechanisms $M: \mathcal{X} \rightarrow \mathbb{D}\mathcal{Y}$ for which $M \downarrow \{x_1, x_2\} \equiv T$.*

Proof: Follows from the construction given in the proof of Thm. A4. ■

And now we show that non-triviality is sufficient for impossibility.

Lemma A6. *There are no universally ℓ -optimal \mathbf{d}_D -private mechanisms over $n > 2$ inputs for any loss function ℓ for which $\ell \downarrow \{x, x'\}$ is non-trivial for all $x \neq x'$.*

Proof: By Cor. A5, it is sufficient to show that for every kernel mechanism K there exists some pair of inputs x_1, x_2 such that $K \downarrow \{x_1, x_2\} \neq T$. Let K be any \mathbf{d}_D -private kernel mechanism on $n > 2$ inputs with corresponding kernel hyper Δ_K . Then $[\Delta_K]$ has at least 2 posteriors, call them δ^1 and δ^2 . Moreover, each δ^i is a vertex and so the \mathbf{d}_D -privacy constraints hold tightly on $n - 1$ input pairs. Pick any 3 inputs x_1, x_2, x_3 such that $\delta_{x_1}^1 = \alpha \delta_{x_2}^1$ and $\delta_{x_2}^2 = \alpha \delta_{x_1}^2$ (ie. the constraints hold in opposite directions). This means that $K \downarrow \{x_1, x_2\} \equiv T$. We now

show that this cannot be true for both $\{x_1, x_3\}$ and $\{x_2, x_3\}$. The 4 possible constructions for δ^1 and δ^2 restricted to $\{x_1, x_2, x_3\}$ are:

Δ_1	δ^1	δ^2	Δ_2	δ^1	δ^2
x_1	α/k	$1/j$	x_1	α/k	$1/j$
x_2	$1/k$	α/j	x_2	$1/k$	α/j
x_3	$1/k$	α/j	x_3	$1/k$	$1/j$

Δ_3	δ^1	δ^2	Δ_4	δ^1	δ^2
x_1	α/k	$1/j$	x_1	α/k	$1/j$
x_2	$1/k$	α/j	x_2	$1/k$	α/j
x_3	α/k	α/j	x_3	α/k	$1/j$

It is clear that for each possible construction of K there exists a pair x_i, x_j st. $K \downarrow \{x_i, x_j\} \neq T$. Therefore there is no K which can be universally ℓ -optimal for any pairwise non-trivial loss function ℓ and thus no universally ℓ -optimal \mathbf{d}_D -mechanisms exist over the space of $n > 2$ inputs. ■

We now have the following:

Theorem 12. *Let ℓ be a strictly monotone loss function on \mathbf{d} for some metric \mathbf{d} . Then there are no universally ℓ -optimal mechanisms in the privacy type $\mathcal{T}_{\mathcal{X}, \mathbf{d}_D}$ over $n > 2$ inputs.* ■

Proof: Follows from Lem. A3 and Lem. A6. ■

And the final corollary of this section which relies on an intermediate result above:

Corollary 13. *There are no universally ℓ_{bin} -optimal mechanisms of type $\mathcal{T}_{\mathcal{X}, \mathbf{d}_D}$ for $|\mathcal{X}| > 2$.*

Proof: This follows from Lem. A6, noting that $\ell_{\text{bin}} \downarrow \{x, x'\}$ is non-trivial for all $x \neq x'$. ■

D. Results supporting §V

Corollary 17. *The additive capacity for privacy type $\mathcal{T}_{\mathcal{X}, \mathbf{d}}$ where $\mathcal{X} = \{x_1, \dots, x_n\}$ is $1 - c$ where c is the optimal value for the linear optimisation problem described by:*

$$\begin{aligned} &\text{minimise } m_{11} + \dots + m_{nn} \\ &\quad \sum_j m_{ij} = 1 \quad 1 \leq i \leq n \\ &\quad m_{ij} \leq m_{kj} \times e^{d(x_i, x_k)} \quad 1 \leq i, k \leq n \\ &\quad m_{ij} \geq 0 \quad 0 \leq i, j \leq n \end{aligned}$$

Proof: Note that Thm. 5 and Def. 8 says that the additive capacity for a mechanism is 1 minus the sum of the column minima. Observe that the constraint set describes mechanisms satisfying a particular privacy type. The result follows provided that we can show that there is such a mechanism whose additive capacity can be computed from the given objective function, i.e. 1 minus the sum of the diagonals. From Thm. 14 we know that there exists a mechanism that optimises the capacity, and from Def. 8 it can be computed from the sum of the column minima. Note that re-ordering of columns is equivalent to refinement, and that refinement preserves membership in a privacy type. Thus if M is an instance of a mechanism exhibiting optimal additive capacity within a given type, we first produce a refinement M' by summing together any columns labelled y, y' for which the column minima $\min_i M_{iy}$ and $\min_i M_{iy'}$ occur for the same i ,

and noting that the additive capacity of M' is the same as for M , with M' possible having fewer columns. However we note that M' 's column minima all occur for *different* values of i . We then re-order the columns of M' , possibly adding in non-negative columns so that the column minima of M'' is equal to the sum of the diagonals. The result follows since M'' 's additive capacity is the same as that of M and is therefore optimal for the whole privacy type. ■

E. Results supporting §VI

Lemma 19. *The (infinite/truncated) Geometric mechanism is a kernel mechanism of privacy type $\mathcal{T}_{\mathbb{N}, \mathbf{d}_2}$.*

Proof: Denote by G, G_t the geometric and truncated geometric mechanisms respectively. The \mathbf{d} -privacy constraints on both G and G_t hold with equality (by construction) and therefore on the posteriors of $[v \triangleright G]$ and $[v \triangleright G_t]$ (by Lem. 2) – which are the same hyper, call it Δ_G [24]. We have now that Δ_G is a vertex hyper. Since G_t is invertible [24] its columns are linearly independent and therefore so are the posteriors of $[v \triangleright G_t] = \Delta_G$ (by linearity). And so Δ_G is a kernel hyper corresponding to both G and G_t . ■

Lemma 21. *The Random Response mechanism is a kernel mechanism in the privacy type $\mathcal{T}_{\mathcal{X}, \mathbf{d}_D}$.*

Proof: It is easy to see (by construction) that R is \mathbf{d}_D -private, since every column has either $R_{x,y} = R_{x',y}$ or $R_{x,y} = \alpha R_{x',y}$ for any $x, x' \in \mathcal{X}$. To show it is a kernel mechanism we need to show that the hyper $\Delta_R = [v \triangleright R]$ has linearly independent inners, their convex hull contains the uniform distribution, and they are vertices in the space of \mathbf{d}_D -private hypers. Observe that R is doubly stochastic and thus the inners of Δ_R are exactly the columns of R . Observe also that each column of R has $n - 1$ constraints holding tightly, and thus we have that R is a vertex mechanism. Now writing R as

$$R = \frac{1}{k} \begin{pmatrix} \alpha & 1 & 1 & \dots & 1 \\ 1 & \alpha & 1 & \dots & 1 \\ 1 & 1 & \alpha & \dots & 1 \\ & & & \dots & \\ 1 & 1 & 1 & \dots & \alpha \end{pmatrix}$$

we can perform basic row operations, subtracting row $k - 1$ from row k for rows n down to 2 to yield:

$$R' = \frac{1}{k} \begin{pmatrix} \alpha & 1 & 1 & \dots & 1 \\ 1 - \alpha & \alpha - 1 & 0 & \dots & 0 \\ 0 & 1 - \alpha & \alpha - 1 & \dots & 0 \\ & & & \dots & \\ 0 & 0 & 0 & \dots & \alpha - 1 \end{pmatrix}$$

Noting that $\det R = \det R'$, we compute:

$$\begin{aligned} \det R' &= \frac{1}{k^n} \left(\alpha(\alpha - 1)^{n-1} - 1(1 - \alpha)(\alpha - 1)^{n-2} + 1(1 - \alpha)^2(\alpha - 1)^{n-3} + \dots + (-1)^{n-1}(1 - \alpha)^{n-1} \right) \\ &= \frac{1}{k^n} \left((\alpha(\alpha - 1)^{n-1} + (\alpha - 1)^{n-1} + (\alpha - 1)^{n-1} + \dots + (\alpha - 1)^{n-1}) \right) \\ &= \frac{1}{k^n} (\alpha - 1)^{n-1} (\alpha + (n - 1)) \end{aligned}$$

which is non-zero for $\alpha \in (0, 1)$. Thus, except for the case where $\alpha = 1$ we have that the randomised response matrix is invertible, and so its columns are linearly independent. And therefore the inners of $\Delta_R = [v \triangleright R]$ are likewise linearly independent. Finally, it is easy to check that we can write the uniform distribution as the following convex combination of columns of R :

$$\frac{1}{n} R_{(-,1)} + \frac{1}{n} R_{(-,2)} + \dots + \frac{1}{n} R_{(-,n)}$$

where $R_{(-,i)}$ denotes the i th column of R . Therefore this also holds for the inners of Δ_R . Thus R is a kernel mechanism as required. ■