

A criterion for privacy protection in data collection and its attainment via randomized response procedures*

Jichong Chai

Department of Statistics, George Washington University, Washington, DC 20052, U.S.A.
e-mail: cjc1990ong@gwmail.gwu.edu

and

Tapan K. Nayak

*Center for Statistical Research and Methodology,
U.S. Census Bureau, Washington, DC 20233, U.S.A. and
Department of Statistics, George Washington University, Washington, DC 20052, U.S.A.*
e-mail: tapan@gwu.edu

Abstract: Randomized response (RR) methods have long been suggested for protecting respondents' privacy in statistical surveys. However, how to set and achieve privacy protection goals have received little attention. We give a full development and analysis of the view that a privacy mechanism should ensure that no intruder would gain much new information about any respondent from his response. Formally, we say that a privacy breach occurs when an intruder's prior and posterior probabilities about a property of a respondent, denoted p and p_* , respectively, satisfy $p_* < h_l(p)$ or $p_* > h_u(p)$, where h_l and h_u are two given functions. An RR procedure protects privacy if it does not permit any privacy breach. We explore effects of (h_l, h_u) on the resultant privacy demand, and prove that it is precisely attainable only for certain (h_l, h_u) . This result is used to define a canonical strict privacy protection criterion, and give practical guidance on the choice of (h_l, h_u) . Then, we characterize all privacy satisfying RR procedures and compare their effects on data utility using sufficiency of experiments and identify the class of all admissible procedures. Finally, we establish an optimality property of a commonly used RR method.

MSC 2010 subject classifications: Primary 62D05; secondary 62B15.

Keywords and phrases: admissibility, Bayes factor, data utility, privacy breach, sufficiency of experiments, transition probability matrix.

Received February 2018.

1. Introduction

In recent years, businesses, organizations and government agencies have been gathering increasingly vast amounts of data from surveys, commercial transactions, on-line searches and postings, medical records and other sources, and heavily using data analytics in making business and policy decisions. Simultaneously, concerns about privacy and data confidentiality have been increasing

*The views expressed in this article are those of the authors and not necessarily those of the U.S. Census Bureau.

substantially. Protecting privacy and personal information is essential for legal reasons and for upholding public trust and support. Several books, e.g., Willenborg and de Waal (2001), Aggarwal and Yu (2008), Hundepool et al. (2012) and Torra (2017), and many papers discuss various privacy and confidentiality protection methods such as grouping, data swapping, cell suppression, imputation and response randomization.

Privacy violations occur in many forms depending on data type, privacy desires and intruders' knowledge and behavior. Thus, various privacy concepts and measures have appeared in the literature, including identity disclosure, differential privacy, k -anonymity and l -diversity (see Chen et al., 2009). Fung et al. (2010) present a systematic review of different approaches. However, as Kifer and Lin (2012) noted, most privacy measures are developed intuitively and can lead us astray, and thus one should use privacy criteria that are logically sound and practical. Evfimievski et al. (2003) introduced one such criterion, called ρ_1 -to- ρ_2 privacy, in the context of randomized response (RR) surveys of categorical variables. Nayak et al. (2015) proposed a similar criterion, called β -factor privacy. The main objectives of this paper are to present some new perspectives on these two criteria and develop and explore the underlying ideas in full generality. Interestingly, we find that any privacy specification amounts to putting an upper bound on all Bayes factors. Thus, privacy needs should be assessed most appropriately in terms of Bayes factors. We obtain a complete characterization of all RR procedures that satisfy any specified privacy criterion. Moreover, we compare all privacy preserving procedures by data utility and identify the admissible procedures.

To describe the context and concepts, we consider a categorical survey variable (or a cross-classification of several variables) X with the set of possible categories $\mathcal{S}_X = \{c_1, \dots, c_k\}$. Let $\pi_i, i = 1, \dots, k$, denote the population level relative frequencies of c_1, \dots, c_k , which are unknown. We collect data to estimate $\pi = (\pi_1, \dots, \pi_k)'$ and make other inferences about π . To protect privacy, an RR survey asks each respondent to use a given random mechanism to generate and report a perturbed version Z of his/her true value of X . We refer to Warner (1965), Chaudhuri and Mukerjee (1988), Chaudhuri (2010) and Nayak et al. (2016) for review of RR theory and additional references. Denote the output space by $\mathcal{S}_Z = \{d_1, \dots, d_m\}$. The transition probabilities $p_{ij} = P(Z = d_i | X = c_j), i = 1, \dots, m, j = 1, \dots, k$, are prespecified and embedded in the randomization device. The matrix $P = ((p_{ij}))$, called the transition probability matrix (TPM), determines all statistical properties of the RR mechanism, and designing an RR survey essentially reduces to choosing P . Thus, we shall identify an RR procedure by its underlying P . We shall require that each row of P contains at least one nonzero element to define m and \mathcal{S}_Z unambiguously. Note that if the i th row of P is zero, then $P(Z = d_i)$ is always zero and d_i is irrelevant. The columns of P are also called channel distributions; see Duchi et al. (2018). Note that the sample spaces \mathcal{S}_X and \mathcal{S}_Z of X and Z , respectively, need not be the same, or even have the same cardinality. For example, in the RAPPOR algorithm of Erlingsson et al. (2014), $m = 2^k$.

Clearly, an RR survey generates data on Z (and not on X). Under simple

random sampling, the distribution of Z is determined by

$$\lambda = P\pi. \quad (1.1)$$

Thus, the data on Z can be used to estimate λ . To estimate π , essentially one would need to use (1.1) and an estimate of λ , say $\hat{\lambda}$. If $k = m$ and P is non-singular, π is estimated using $\hat{\pi} = P^{-1}\hat{\lambda}$, see e.g., Chaudhuri and Mukerjee (1988). If $m < k$, or more generally if the columns of P are not linearly independent, the model for Z is not identifiable with respect to π and hence π is not estimable. Thus, inference and data utility considerations suggest to use $m \geq k$ and P with rank k . In a parametric model, if π is a function of fewer parameters, identifiability might hold even when $m < k$. However, identifiability with respect to π ensures that the data set can be analyzed using different models for various (possibly unforeseen) purposes.

In the preceding framework, Evfimievski et al. (2003) defined ρ_1 -to- ρ_2 privacy, taking a Bayesian view. For a target respondent B , suppose an intruder R 's prior probability of $X = c_j$ is α_j , and let $\alpha = (\alpha_1, \dots, \alpha_k)'$. Note that an intruder's prior α about a target may be quite different from π . For a given prior α , $P_\alpha(Z = d_i) = \sum_{l=1}^k \alpha_l p_{il}$ and the posterior probability of $X = c_j$ given B 's response $Z = d_i$, is

$$P_\alpha(X = c_j | Z = d_i) = \frac{P_\alpha(X = c_j, Z = d_i)}{P_\alpha(Z = d_i)} = \frac{\alpha_j p_{ij}}{\sum_{l=1}^k \alpha_l p_{il}}. \quad (1.2)$$

Also, R 's prior and posterior probabilities of any $Q \subseteq \mathcal{S}_X = \{c_1, \dots, c_k\}$ are:

$$P_\alpha(X \in Q) = \sum_{j:c_j \in Q} \alpha_j \quad \text{and} \quad P_\alpha(X \in Q | Z = d_i) = \sum_{j:c_j \in Q} P_\alpha(X = c_j | Z = d_i).$$

For brevity, we shall denote $P_\alpha(X \in Q)$ by $P_\alpha(Q)$ and $P_\alpha(X \in Q | Z = d_i)$ by $P_\alpha(Q | d_i)$.

Definition 1.1. (Evfimievski et al., 2003) Let $0 < \rho_1 < \rho_2 < 1$ be two numbers. (a) An RR procedure is said to permit an upward ρ_1 -to- ρ_2 privacy breach with respect to $Q \subseteq \mathcal{S}_X$ and a prior distribution α if for some $1 \leq i \leq m$ with $P_\alpha(Z = d_i) > 0$,

$$P_\alpha(Q) < \rho_1 \quad \text{and} \quad P_\alpha(Q | d_i) > \rho_2. \quad (1.3)$$

Similarly, a procedure is said to admit a downward ρ_2 -to- ρ_1 privacy breach if $P_\alpha(Q) > \rho_2$ and $P_\alpha(Q | d_i) < \rho_1$ for some d_i with $P_\alpha(Z = d_i) > 0$.

(b) An RR procedure is said to provide ρ_1 -to- ρ_2 privacy protection if it does not permit an upward ρ_1 -to- ρ_2 privacy breach or a downward ρ_2 -to- ρ_1 privacy breach with respect to any Q and any prior α .

Definition 1.2. (Nayak et al., 2015) For a given $\beta > 1$, an RR procedure admits a β -factor privacy breach, with respect to $Q \subseteq \mathcal{S}_X$ and a prior α if $P_\alpha(Q) > 0$ and

$$\frac{P_\alpha(Q | d_i)}{P_\alpha(Q)} > \beta \quad \text{or} \quad \frac{P_\alpha(Q | d_i)}{P_\alpha(Q)} < \frac{1}{\beta} \quad (1.4)$$

for some d_i such that $P_\alpha(Z = d_i) > 0$.

An RR procedure provides β -factor privacy if it does not allow a β -factor breach with respect to any Q and any α .

The above two criteria are very strong, as they require no privacy breach for any d_i, Q and α . Thus, no answer (d_i) of a respondent B would give “much” new information to any intruder R (characterized by α) about any property (Q) of B with respect to X . Evfimievski et al. (2004) introduced a similar concept of privacy breach in privacy preserving association rule mining. In practice, values of (ρ_1, ρ_2) and β should be chosen based on the sensitivity of X and privacy concerns. Here, the β -factor privacy is simpler as it requires us to specify only one number (β). Interestingly, the strict privacy requirements of the two criteria are achievable, as summarized below.

Definition 1.3. (Nayak et al., 2015) The i th row parity of P is defined as

$$\eta_i(P) = \max \left\{ \frac{p_{ij}}{p_{il}} \mid j, l = 1, \dots, k \right\} = \frac{\max_j \{p_{ij}\}}{\min_j \{p_{ij}\}}, \quad (1.5)$$

with the convention $0/0 = 1$ and $a/0 = \infty$ for any $a > 0$.

Furthermore, the parity of P is defined as $\eta(P) = \max_i \{\eta_i(P)\}$.

Theorem 1.1. (Evfimievski et al., 2003) A sufficient condition for an RR procedure with transition probability matrix P to guarantee ρ_1 -to- ρ_2 privacy is:

$$\eta(P) \leq \frac{\rho_2(1 - \rho_1)}{\rho_1(1 - \rho_2)}. \quad (1.6)$$

Theorem 1.2. (Nayak et al., 2015) An RR procedure guarantees β -factor privacy if and only if $\eta(P) \leq \beta$.

We shall see later that (1.6) is also a necessary condition for P to provide ρ_1 -to- ρ_2 privacy. We should mention that Boreale and Paolini’s (2015) concept of “worst-case breach” is essentially the same as β -factor breach. They also proved a version of Theorem 1.2. The concept of parity is very similar to γ -amplification of Evfimievski et al. (2003). Clearly, $\eta(P) \geq 1$ and it is finite if and only if all elements of P are positive. Also, for any given η_0 , it is possible to construct P with parity η_0 ; see Evfimievski et al. (2003) and Agrawal et al. (2009). In particular, for $m \geq k$, one P with $\eta(P) = \eta_0$ is obtained by taking $p_{ii} = \eta_0 / [\eta_0 + m - 1]$, $i = 1, \dots, k$, and $p_{ij} = 1 / [\eta_0 + m - 1]$ for all $i \neq j$.

The rest of the paper is organized as follows. In Section 2, we present some new perspectives on ρ_1 -to- ρ_2 and β -factor privacy, including a geometric view and equivalency with ϵ -differentially local privacy, and then propose a general privacy criterion (in Definition 2.2) that covers Definitions 1.1 and 1.2 as special cases. Essentially, we pursue the spirit of ρ_1 -to- ρ_2 and β -factor privacy to the fullest extent and permit any (reasonable) privacy breach criterion. In Section 3, we explore implications and practicality of the general criterion. We develop a canonical form of the general criterion and characterize all RR procedures that provide required privacy. In Theorem 3.1, we prove that P satisfies a specified

privacy demand if and only if $\eta(P)$ is appropriately small. In Section 4, we compare data utility of all privacy satisfying P . Employing Blackwell's (1951, 1953) concept of sufficiency of experiments, which is agnostic about inferential goals and loss functions, we characterize the class of all admissible privacy preserving procedures. We also prove a particular optimality property of a simple RR procedure. We note some concluding remarks in Section 5.

2. A general criterion

To motivate a general criterion, we shall first discuss some logical and practical features of Definitions 1.1 and 1.2. The ρ_1 -to- ρ_2 and β -factor privacy criteria are very strong, but it should be noted that those are applicable only when an intruder R knows his/her target B 's value of Z . Typically, this happens at data collection time, with R being the data collector. In commercial data mining context, Agrawal et al. (2009) refer to this as business-to-customer (or B2C) privacy. Definitions 1.1 and 1.2 are not applicable if R gets access only to an anonymized version of the original data set, where B 's records cannot be ascertained with certainty. In other words, ρ_1 -to- ρ_2 and β -factor privacy criteria presumes disclosure of B 's identity to an intruder.

Related to the preceding point, we also want to mention that while privacy and confidentiality have often been used synonymously, those should be distinguished due to some important differences (see Nayak et al., 2015). In legal terms, privacy is a person's right to freedom from intrusion into his/her information. Privacy emerges as a desire to share no or only obscured information with a data collector. Thus, privacy protection should occur at the time of data collection. In contrast, confidentiality is an obligation to prevent unauthorized access to private information. People often give their information trusting that their data will be used by researchers and policy makers only to learn about the population as a whole and not about any individual. Privacy applies to individuals whereas confidentiality applies to the data, which may be addressed after data collection. One important technical (and practical) implication is that one may examine the whole data set for choosing a suitable method for confidentiality protection. In contrast, methods for privacy protection need to be selected before data collection. Consequently, some concepts, such as k -anonymity and l -diversity, and related methods are applicable for protecting confidentiality but not privacy.

Remark 2.1. *A referee suggested us to relate this paper to research on PRAM (the post-randomization method), introduced by Gouweleew et al. (1998). Many papers, e.g., Van den Hout and Van der Heijden (2002), Van den Hout and Elamir (2006) and Cruyff et al. (2008), discuss properties, variations and applications of the method. PRAM and RR are closely related as both methods randomize true responses using a transition probability matrix P . However, PRAM and RR are not equivalent (see Nayak et al., 2015 and Nayak and Adeshiyan, 2016). RR is for privacy protection whereas PRAM aims to protect confidentiality. Consequently, one important distinction is that P is fixed in RR, but in*

PRAM, it can be chosen based on the observed (unperturbed) data set. In particular, in invariant (or unbiased) PRAM, P must depend on the data. Such data dependency of P affects statistical inferences, which was noted and investigated by Nayak and Adeshiyan (2016).

Privacy and confidentiality concerns for are different. So, their measures and protection goals are also different. One of the top confidentiality breach concerns is identity disclosure, which occurs when the records of a survey unit is correctly identified in released data, by matching the values of some of the variables that are easily available from other sources. Recently, Nayak et al. (2018) developed a PRAM procedure that strictly controls identification risks at a very little loss of data utility. In summary, PRAM a useful tool for confidentiality protection, but this paper, which deals with privacy protection, is not directly linked to PRAM.

As we discuss next, logically ρ_1 -to- ρ_2 and β -factor privacy directly address the core of privacy concern, which is: how much *information* an intruder might gain about a respondent from his/her response (possibly perturbed)? One compelling view of information, as Basu (1988) articulated, is: “Information is what information does. It changes opinion.” Furthermore, opinion can be expressed precisely only using subjective probability. An intruder’s prior and posterior probabilities describe respectively his/her initial and revised opinion, after learning a respondent’s reported value. These constitute a strong argument that privacy should be discussed in terms of intruders’ prior and posterior probabilities (instead of technical information measures, e.g., mutual information and f -divergence, that were developed in other contexts). Definitions 1.1 and 1.2 coincide with the above view and are thus highly relevant to privacy considerations.

The changing of a prior to posterior occurs only through the likelihood function, and the change is small if the likelihood function is relatively flat. In our setting, for response d_i , the likelihoods for c_1, \dots, c_k are $p_{ij}, j = 1, \dots, k$, and they are fairly close to each other when $\eta_i(P)$ is small. Consequently, the likelihood functions for all possible responses are fairly flat if and only if $\eta(P)$ is fairly small. This comes out precisely in Theorems 1.1 and 1.2.

We now mention a connection to the following concept (see, Duchi et al., 2018) of differential local privacy.

Definition 2.1. An RR method provides ϵ -differentially local privacy (ϵ -DLP), for $\epsilon > 0$, if

$$\max \left\{ \frac{P(Z \in S|X = c_j)}{P(Z \in S|X = c_l)} \mid S \subseteq \mathcal{S}_Z, 1 \leq j, l \leq k \right\} \leq \exp(\epsilon). \quad (2.1)$$

It can be seen that (2.1) is equivalent to $\eta(P) \leq \exp(\epsilon)$. So, in view of Theorem 1.2, ϵ -DLP and β -factor privacy are equivalent, with $\beta = \exp(\epsilon)$. An equivalency of ϵ -DLP and ρ_1 -to- ρ_2 privacy can be observed similarly. Clearly, the thinking behind Definitions 1.1, 1.2 and 2.1 are different, but mathematically, they are equivalent as each one corresponds to an upper bound on $\eta(P)$. Figure 1




FIG 1. ρ_1 -to- ρ_2 and β -factor privacy breach regions.

gives a helpful geometrical perspective of ρ_1 -to- ρ_2 and β -factor privacy. The two shaded rectangles represent the privacy breach region (PBR) of ρ_1 -to- ρ_2 privacy, as any (prior, posterior) pair, to be denoted generically by (p, p_*) , falling in this region signifies a privacy breach. The two shaded triangles constitute the PBR of β -factor privacy. In practice, visual inspection of various PBRs might help to choose the parameter values, e.g., (ρ_1, ρ_2) or β , of a privacy criterion, and also to compare different privacy guarantees. Naturally, a larger PBR implies a stronger privacy guarantee. Among two PBRs in Figure 1, none is a subset of the other one, but as the β -factor PBR has a larger area and covers most of the other PBR, one might reasonably consider it stronger. As such, two overlapping PBRs, as in Figure 1, are not comparable, but we shall see in Section 3 that privacy demands of any two PBRs can be compared meaningfully.

One main goal of this paper is to explore the central idea of ρ_1 -to- ρ_2 and β -factor privacy in full generality. Thus, we now consider a general privacy breach region W , as shown by the shaded region in Figure 2, and require that no (prior, posterior) pair must fall in W . So, the unshaded part (W^c) is the privacy holding region. Describing the down and up privacy breach boundaries of W with two functions h_l and h_u , we introduce the following.

Definition 2.2. Let h_l and h_u be two functions from $[0, 1]$ to $[0, 1]$ such that $0 \leq h_l(a) \leq a \leq h_u(a) \leq 1$ for all $0 \leq a \leq 1$. An RR procedure is said to provide strict information privacy (SIP) with respect to h_l and h_u , to be abbreviated (h_l, h_u) -SIP, if

$$h_l(P_\alpha(Q)) \leq P_\alpha(Q|d_i) \leq h_u(P_\alpha(Q)) \quad (2.2)$$

for $i = 1, \dots, m$ and all $\alpha, Q \subseteq \mathcal{S}_X$.




FIG 2. A general privacy breach region.

Clearly, the general idea is that for privacy protection, if the prior probability of an event is p , then its posterior probability must be between $h_l(p)$ and $h_u(p)$. Obviously, this covers Definitions 1.1 and 1.2 as special cases. Mathematically, we do not need to put additional conditions on h_l and h_u , but intuitively, they should be nondecreasing. (We shall see in the sequel that all precise PBRs satisfy this condition.) Definition 2.2 specifies a *privacy demand* with its PBR $W[h_l, h_u] = \{(p, p_*), 0 \leq p, p_* \leq 1 : p_* < h_l(p) \text{ or } p_* > h_u(p)\}$. On the other hand, the *privacy provided* by any RR procedure can be described by its PBR as defined next.

Definition 2.3. We define the PBR of any RR procedure P as the collection of all non-attainable (prior, posterior) pairs under P , and denote it by W_P . Thus, W_P is the complement (with respect to the unit square) of P 's privacy holding region: $\{(p, p_*), 0 \leq p, p_* \leq 1 : P_\alpha(Q) = p \text{ and } P_\alpha(Q|d_i) = p_* \text{ for some } d_i, \alpha \text{ and } Q \subseteq \mathcal{S}_X\}$.

Definition 2.4. We shall call a general privacy breach region W precise if there exists an RR procedure P such that $W_P = W$.

The preceding two definitions will be useful to comparing and matching privacy demand with privacy provided by different procedures. Clearly, an RR procedure P provides (h_l, h_u) -SIP if and only if $W[h_l, h_u] \subseteq W_P$. However, if $W[h_l, h_u]$ is not precise, to guarantee (h_l, h_u) -SIP one must use an RR procedure P for which W_P is strictly larger than $W[h_l, h_u]$, and in such cases, we should report W_P , the PBR of the procedure actually used, to communicate the privacy guarantee precisely and maximally. This also implies that to determine privacy requirement we should think only in terms of precise PBRs. These

observations raise some natural questions, such as: What are the precise PBRs? Which procedures satisfy a given precise PBR? For given h_l and h_u , is there a minimal W_P satisfying $W[h_l, h_u] \subseteq W_P$? We answer these questions in the next section.

3. Characterization of strict information privacy

We begin this section with some analytic simplifications of the (h_l, h_u) -SIP criterion. First, note that $P_\alpha(Q) = 0$ implies $P_\alpha(Q|d_i) = 0$ and $P_\alpha(Q) = 1$ implies $P_\alpha(Q|d_i) = 1$, for all d_i . So, (2.2) holds automatically when $P_\alpha(Q)$ is 0 or 1, and to establish (h_l, h_u) -SIP, we need to verify (2.2) only for all $\alpha, Q \subseteq \mathcal{S}_X$ such that $0 < P_\alpha(Q) < 1$. Second, observe that the first \leq in (2.2) is equivalent to $1 - h_l(P_\alpha(Q)) \geq 1 - P_\alpha(Q|d_i)$ or $P_\alpha(Q^c|d_i) \leq 1 - h_l(1 - P_\alpha(Q^c))$. So, the first \leq in (2.2) holds for all $Q \subseteq \mathcal{S}_X$ if and only if $P_\alpha(Q|d_i) \leq 1 - h_l(1 - P_\alpha(Q))$ for all $Q \subseteq \mathcal{S}_X$, i.e., the condition for downward privacy breach is equivalent to an upward privacy breach criterion. (Evfimievski et al., 2003) made a similar observation for ρ_1 -to- ρ_2 privacy.) Combining the two upward breach conditions and defining $[h_l \star h_u](a) = \min\{h_u(a), 1 - h_l(1 - a)\}$, for $0 \leq a \leq 1$, we obtain the following:

Lemma 3.1. *Let h_l and h_u be as in Definition 2.2 and $[h_l \star h_u]$ be defined as above. Then, an RR procedure P provides (h_l, h_u) -SIP if and only if*

$$P_\alpha(Q|d_i) \leq [h_l \star h_u](P_\alpha(Q)) \quad (3.1)$$

for all $i = 1, \dots, m$ and all α and $Q \subseteq \mathcal{S}_X$ such that $0 < P_\alpha(Q) < 1$.

The conditions $0 \leq h_l(a) \leq a \leq h_u(a) \leq 1$ imply that $a \leq [h_l \star h_u](a) \leq 1$ for all $0 \leq a \leq 1$. In view of Lemma 3.1 and preceding discussions, we may define a privacy criterion more succinctly only in terms of upward breaches as follows.

Definition 3.1. *Let $h : [0, 1] \rightarrow [0, 1]$ be a function satisfying $a \leq h(a) \leq 1$ for all $0 \leq a \leq 1$. An RR procedure is said to provide canonical strict information privacy with respect to h , to be abbreviated h -CSIP, if*

$$P_\alpha(Q|d_i) \leq h(P_\alpha(Q)) \quad (3.2)$$

for $i = 1, \dots, m$ and all α and $Q \subseteq \mathcal{S}_X$ such that $0 < P_\alpha(Q) < 1$.

It can be seen that h -CSIP also provides the downward privacy guarantee that $P_\alpha(Q|d_i) \geq \tilde{h}(P_\alpha(Q))$ for $i = 1, \dots, m$ and all α and $Q \subseteq \mathcal{S}_X$, where $\tilde{h}(a) = 1 - h(1 - a)$, $0 \leq a \leq 1$. Thus, the upper and lower boundaries of the PBR of h -CSIP are given by h and \tilde{h} , respectively. Lemma 3.1 shows that for any h_l and h_u , (h_l, h_u) -SIP and $[h_l \star h_u]$ -CSIP are equivalent, in the sense that if an RR procedure guarantees one of the two, it must also guarantee the other one. However, the PBR given by some (h_l, h_u) -SIP can be a proper subset of the PBR of the corresponding $[h_l \star h_u]$ -CSIP. This is illustrated in Figure 3, where the PBR of $[h_l \star h_u]$ -CSIP is the PBR of (h_l, h_u) -SIP (shown as the region shaded




FIG 3. Connection between privacy breach regions of (h_l, h_u) -SIP and $[h_l \star h_u]$ -CSIP.

with solid lines) plus the two dotted lined parts A and B. The two PBRs would be identical only when $h_l(a) = 1 - h_u(1 - a), 0 \leq a \leq 1$. While Definition 3.1 is technically most concise, in real applications, it might be more convenient to specify h_l and h_u , defining lower and upper privacy breaches, and then take $h = [h_l \star h_u]$.

Subsequently, we shall explore only h -CSIP, because it is the analytical crux of any privacy criterion as seen above. For any given h , define

$$B(h) = \inf_{0 < p < 1} \left(\frac{1-p}{p} \right) \left(\frac{h(p)}{1-h(p)} \right), \tag{3.3}$$

where we take $h(p)/[1-h(p)] = \infty$ when $h(p) = 1$. (Alternatively, we can take the infimum over $\{0 < p < 1 : h(p) < 1\}$.) The following results characterize all RR procedures that provide h -CSIP, for any given h .

Lemma 3.2. An RR procedure P provides h -CSIP, with a specified h , if and only if

$$P_\alpha(X = c_j | Z = d_i) \leq h(P_\alpha(X = c_j)) \tag{3.4}$$

for all i, j and α such that $P_\alpha(X = c_j) > 0$ and $P_\alpha(Z = d_i) > 0$.

Lemma 3.3. A necessary and sufficient condition for an RR procedure P to satisfy (3.4) for all i, j and α is that $\eta(P) \leq B(h)$.

Theorem 3.1. For any given h , an RR procedure P provides h -CSIP if and only if $\eta(P) \leq B(h)$.

Interestingly, Lemma 3.2 says that to assure (3.2) for all $Q \subseteq \mathcal{S}_X$, it suffices to verify (3.2) only for $\{X = c_j\}$, i.e., for the atomic events of X . The ‘only if’ part of this lemma is obvious and Theorem 3.1 follows readily from the two

lemmas. The remaining proofs are given in the Appendix. Theorems 1.1 and 1.2 can be obtained from Theorem 3.1 by calculating $B(h)$ for relevant h functions.

The necessary and sufficient condition in Theorem 3.1 depends on h only through $B(h)$ and on P only through its parity $\eta(P)$. Thus, in h -CSIP context, $B(h)$ quantifies the privacy demand of h and $\eta(P)$ is the privacy level of P . We can measure of the privacy demand of any general PBR, with downward and upward breach boundaries h_l and h_u , as $B(h)$, where $h = [h_l \star h_u]$; recall that $[h_l \star h_u](a) = \min\{h_u(a), 1 - h_l(1 - a)\}$, $0 \leq a \leq 1$. Using this measure, we can compare privacy demands of any two PBRs, even when they overlap as in Figure 1. Likewise, we can compare the privacy level of all RR procedures P using parity.

Consider an RR procedure P with $\eta(P) = \gamma > 1$. Then, by Theorem 3.1 and (3.3), P guarantees h -CSIP for all h such that

$$\gamma \leq \left(\frac{1-p}{p}\right) \left(\frac{h(p)}{1-h(p)}\right) \quad \text{for all } 0 < p < 1,$$

or equivalently $h(p) \geq h_{(\gamma)}(p)$, where $h_{(\gamma)}(p)$ is defined as

$$h_{(\gamma)}(p) = \frac{\gamma p}{1 + (\gamma - 1)p}, \quad 0 < p < 1. \quad (3.5)$$

Thus, $h_{(\gamma)}(\cdot)$ is the up breach boundary of any P with parity γ . The corresponding down boundary is $\tilde{h}_{(\gamma)}(p) = 1 - h_{(\gamma)}(1 - p)$. Note that the PBR of P is determined only by its parity. So, all P with a common parity have the same PBR. It also follows that W is a precise PBR if and only if its up and down breach boundaries are $h_{(\gamma)}$ and $\tilde{h}_{(\gamma)}$, respectively, for some $\gamma > 1$. As γ increases, $h_{(\gamma)}(p)$ shifts upward and the PBR gets smaller, as shown in Figure 4.

Let $\mathcal{H} = \{h_{(\gamma)}(\cdot); \gamma > 1\}$, i.e., the class of all function of the form $h_{(\gamma)}(\cdot)$. Then, h -CSIP with all $h \in \mathcal{H}$ represent all precise PBRs, which are most relevant to choosing privacy requirement and communicating privacy guarantee. A logical conclusion is that for strict privacy protection, we should think only in terms of h -CSIP and limit h to \mathcal{H} . A practical meaning of $h_{(\gamma)}$ -CSIP may not be immediate, but as we show next, this amounts to imposing a bound on all Bayes factors.

For given α , the prior odds of Q is $P_\alpha(Q)/[1 - P_\alpha(Q)]$ and its posterior odds given $Z = d_i$ is $P_\alpha(Q|d_i)/[1 - P_\alpha(Q|d_i)]$. Now, take any $\gamma > 1$ and consider the following privacy requirement:

$$\frac{P_\alpha(Q|d_i)}{1 - P_\alpha(Q|d_i)} \frac{1 - P_\alpha(Q)}{P_\alpha(Q)} \leq \gamma \quad (3.6)$$

for all α, Q and d_i such that $0 < P_\alpha(Q) < 1$ and $P_\alpha(Z = d_i) > 0$. The left side of (3.6) is the ratio of posterior odds of Q to its prior odds, or the Bayes factor for testing $X \in Q$ against $X \notin Q$; see Kass and Raftery (1995) for a very informative discussion of Bayes factor. Thus, (3.6) requires all Bayes factors to




FIG 4. Plots of precise up breach boundaries $h_{(\gamma)}$.

be less than or equal to γ . Considering Q^c , it can be seen that (3.6) also implies that all Bayes factors are at least γ^{-1} . In summary, (3.6) requires all Bayes factors to be between γ^{-1} and γ . The above criteria is analogous to β -factor privacy; while (3.6) uses the ratio of posterior and prior odds, β -factor privacy uses the ratio of the two probabilities. In other words, β -factor privacy uses probability scale whereas (3.6) uses odds scale.

By routine algebra, it can be seen that (3.6) is equivalent to

$$P_\alpha(Q|d_i) \leq \frac{\gamma P_\alpha(Q)}{1 + (\gamma - 1)P_\alpha(Q)}. \tag{3.7}$$

Notice that the right side of (3.7), considered as a function of $P_\alpha(Q)$, is the same as the function $h_{(\gamma)}(\cdot)$ defined in (3.5). So, $h_{(\gamma)}$ -CSIP is equivalent to the privacy requirement of (3.6) and γ can be interpreted conveniently as the upper bound on all Bayes factors (and the lower bound is γ^{-1}). It is also seen that any precise PBR corresponds to the privacy requirement of (3.6), with a matching value of γ . Based on previous discussions we reach the following practical conclusions.

(1) While ρ_1 -to- ρ_2 and β -factor privacy and more generally (h_l, h_u) -SIP are intuitively sensible, we should discuss, assess and communicate privacy only in terms of h -CSIP with $h \in \mathcal{H}$, or equivalently in terms of bounds on Bayes factors as in (3.6). Both the graphical representation, as in Figure 4, and the Bayes factor interpretation should be helpful for choosing suitable values of γ in practical applications. Kass and Raftery (1995) recommend to interpret a Bayes factor 20 or larger as strong evidence, which suggests that values around 20 might be suitable for γ in our context.

(2) Satisfying any privacy requirement reduces strictly to using a procedure with a sufficiently small parity, as stated in Theorem 3.1. We can always find a

procedure to provide required privacy, but not uniquely because for any $\gamma > 1$, there exists many P with $\eta(P) = \gamma$. We should compare data utility to choose one procedure among all privacy satisfying procedures. We discuss this in the next section.

(3) Our results also bring out a new interpretation of ϵ -DLP (see Definition 2.1). Recall that an RR procedure P provides ϵ -DLP if and only if $\eta(P) \leq \exp(\epsilon)$. Then, in view of Theorem 3.1 and subsequent discussions, ϵ -DLP is equivalent to $h_{(\gamma)}$ -CSIP, with $\gamma = \exp(\epsilon)$. Thus, ϵ -DLP can be explained by the PBR of the corresponding $h_{(\gamma)}$ -CSIP.

4. Comparison of data utility

In earlier sections, we observed that a randomization procedure P provides strict privacy protection if and only if $\eta(P) \leq \gamma$, where $\gamma > 1$ is determined by the privacy requirement. Recall that $P_{m \times k}$ must be a transition probability matrix (TPM) and each row of P must contain at least one nonzero value. We also argue that no two rows of P should be proportional to each other. The i th row is the (nonparametric) likelihood function when $Z = d_i$. So, from likelihood perspective, if rows i and j are proportional, the statistical information from observing $Z = d_i$ and $Z = d_j$ are the same and the two outcomes (and the corresponding rows) should be merged. Alternatively, two proportional rows can be viewed as obtained from randomly splitting one outcome into two. (This is analogous to irrelevantly splitting one choice into two in discrete choice analyses; e.g., splitting “bus” into “blue bus” and “green bus” in mode of transportation choice.) Also, if proportional rows are allowed, then \mathcal{S}_Z and m cannot be defined uniquely.

With the natural constraints discussed above, the class of all privacy preserving procedures, at a desired level γ , is:

$$\mathcal{C}_\gamma = \{P_{m \times k} : m \geq 2, \eta(P) \leq \gamma \text{ and } P \text{ has no proportional rows}\}.$$

As we noted earlier, one may also impose $m \geq k$ and $\text{rank}(P) = k$, for model identifiability. However, these are not needed for our results. Intuitively, we should compare data utility to select a procedure from \mathcal{C}_γ for application. However, “data utility” is difficult to define and measure as the data may be used and analyzed in different ways and for various purposes. It may not even be possible to anticipate all future usage of the data at the time of the survey. Recognizing this, we shall first discuss some admissibility results using Blackwell’s (1951, 1953) notion of *sufficiency of experiments*, which is agnostic about inferential goals and loss functions.

4.1. Admissibility

Adopting Blackwell’s (1951, 1953) criterion to our context, we introduce the following:

Definition 4.1. For two randomization procedures $A_{r \times k}$ and $P_{m \times k}$, we say that P is at least as informative (or good) as A , to be denoted $P \succeq A$, if there exists a transition probability matrix $C_{r \times m}$ such that $A = CP$. In this case, P is also said to be sufficient for A .

If $P \succeq A$ and also $A \succeq P$, then A and P are equivalent and will be denoted $P \sim A$. We say that P is better than A and write $P \succ A$ if $P \succeq A$ but A is not sufficient for P , i.e., $A \not\succeq P$.

Definition 4.2. A randomization procedure $P \in \mathcal{C}_\gamma$ is said to be inadmissible within \mathcal{C}_γ if there exists $A \in \mathcal{C}_\gamma$ such that $A \succ P$. Otherwise, P is called admissible.

It is easy to see that if C and P are TPMs, then $A = CP$ is also a TPM. The intuitive idea behind Definitions 4.1 and 4.2 is that if $A = CP$, then the procedure A is equivalent to further randomizing (by C) the output of P , and because of the additional randomization, A cannot be more informative than P . Mathematically, it follows that if P is sufficient for A , then for any inference problem with a given loss function and any inference rule δ based the data from A , there exists a rule δ_* based on P such that the risk of δ_* is never larger than the risk of δ . Naturally, one should use only admissible procedures. In privacy literature, Blackwell’s criterion has been used by Kairouz et al. (2016b).

Remark 4.1. The restriction that our TPMs must not contain proportional rows can be further justified as follows. Consider a procedure $A_{m \times k}$ and suppose its first two rows, denoted \vec{a}_1 and \vec{a}_2 are proportional and $\vec{a}_1 = \delta(\vec{a}_1 + \vec{a}_2)$, $0 < \delta < 1$. Construct $P_{(m-1) \times k}^*$ by merging the first two rows of A . Then, A and P^* are equivalent, as $A = CP^*$ and $P^* = C^*A$ with C and C^* defined as:

$$C = \left(\begin{array}{c|c} \delta & 0 \\ 1 - \delta & 0 \\ \hline 0 & I \end{array} \right) \quad \text{and} \quad C^* = \left(\begin{array}{c|c} 1 & 1 \\ 0 & 0 \\ \hline 0 & I \end{array} \right).$$

We can repeat this process to eliminate all proportional rows and thus obtain a P such that P has no proportional rows and $P \sim A$.

Remark 4.2. Intuitively, permuting the rows of $P_{m \times k}$, i.e., relabeling the elements of \mathcal{S}_Z , should have no effect on either privacy or data utility. Mathematically, this holds easily. Specifically, it can be seen that if $C_{m \times m}$ is a permutation matrix and $A = CP$, then (i) $\eta(A) = \eta(P)$ and (ii) $A \sim P$ (as C^{-1} is a also a permutation matrix and hence TPM). We also have the following result, whose proof is given in the Appendix.

Theorem 4.1. Two procedures $A_{m \times k}, P_{r \times k} \in \mathcal{C}_\gamma$ are equivalent if and only if $m = r$ and $A = CP$, where C is a permutation matrix.

Lemma 4.1. Suppose $C_{m \times r}$ and $P_{r \times k}$ are two TPMs, $\eta(P) = \gamma$ and let $A_{m \times k} = CP$. Then, $\eta(A) \leq \gamma$.

Proof. Note that $\eta(P) = \gamma$ implies that $p_{uj} \leq \gamma p_{ul}$ for all u, j and l . So, for all

i, j and l ,

$$a_{ij} = \sum_{u=1}^r c_{iu} p_{uj} \leq \sum_{u=1}^r c_{iu} (\gamma p_{ul}) = \gamma a_{il} \quad (4.1)$$

and thus $\eta_i(A) \leq \gamma$ for $i = 1, \dots, m$, and consequently $\eta(A) \leq \gamma$. \square

This result is intuitive: further randomization should not reduce privacy (by increasing parity). It also exhibits a trade-off between privacy and data utility: if P is at least as informative as A , in the sense of $P \succeq A$, then A provides at least as much privacy (by parity measure) as P .

Next, let \mathcal{C}_γ^0 denote all P in \mathcal{C}_γ satisfying the following two conditions:

C1: $\eta_i(P) = \gamma$ for all i

C2: Each row of P contains exactly two distinct values.

We shall prove that a randomization procedure P is admissible within \mathcal{C}_γ if and only if $P \in \mathcal{C}_\gamma^0$. We organize this result in several parts.

Theorem 4.2. *Any randomization procedure $A \in \mathcal{C}_\gamma^0$ is admissible within \mathcal{C}_γ .*

Proof. Take any $A \in \mathcal{C}_\gamma^0$. We shall prove that if any $P \in \mathcal{C}_\gamma$ is sufficient for A , then A must be equivalent to P . Suppose there exist $P_{r \times k} \in \mathcal{C}_\gamma$ and a TPM $C_{m \times r}$ such that $A = CP$. Each row of C must contain at least one nonzero element, as A does not have any zero row. We shall see that $c_{iu} \neq 0$ implies that the u th row of P , denoted \vec{p}_u , is proportional to \vec{a}_i , the i th row of A . For each i , as $\eta_i(A) = \gamma$, by C1, there exist j and l such that $a_{ij} = \gamma a_{il}$. For such a_{ij} and a_{il} , equality holds in (4.1) and since $c_{iu} \neq 0$, we must have $p_{uj} = \gamma p_{ul}$. This holds for all j and l such that $a_{ij} = \gamma a_{il}$. Since \vec{a}_i contains exactly two distinct (nonzero) values, by C2, considering all pairs (a_{ij}, a_{il}) with $a_{ij} = \gamma a_{il}$, it is seen that $\vec{p}_u \propto \vec{a}_i$.

The preceding result implies that each row of C has exactly one nonzero entry; otherwise, P will have proportional rows. Then, if $m < r$, C must have some zero columns hence would not be a TPM. Also, if $m > r$, at least two rows of C must be proportional, and the corresponding rows of A are also proportional, which is a contradiction. So, we must have $m = r$ and C must be a permutation matrix, to be a TPM, and thus $A \sim P$. \square

We should note that in the preceding proof we not only showed that $A \sim P$ but also that P must be a permutation of the rows of A . Consequently, $P \in \mathcal{C}_\gamma^0$, as $A \in \mathcal{C}_\gamma^0$, and we have following.

Corollary 4.1. *If $A \in \mathcal{C}_\gamma^0$, then no $P \in (\mathcal{C}_\gamma \setminus \mathcal{C}_\gamma^0)$ can be equivalent to A . Stated another way, if $A \in \mathcal{C}_\gamma^0$, $P \in (\mathcal{C}_\gamma \setminus \mathcal{C}_\gamma^0)$ and $A \succeq P$, then $A \succ P$.*

Lemma 4.2. *Suppose $A_{m \times k} \in \mathcal{C}_\gamma$ and $1 < \eta_i(A) < \gamma$ for some i . Then, A is inadmissible.*

Proof. Suppose, without loss of generality, $1 < \eta_1(A) = a_{11}/a_{12} < \gamma$. Then, there exists a row i such that $a_{i1} < a_{i2}$ because each column of A adds to 1. For brevity suppose that $a_{21} < a_{22}$. Construct $P_{m \times k}$ as follows: $\vec{p}_1 = \vec{a}_1 + (1 - \xi)\vec{a}_2$, $\vec{p}_2 = \xi\vec{a}_2$, where $\xi \geq 1$ is a constant (to be chosen suitably) and

$\vec{p}_i = \vec{a}_i, i = 3, \dots, m$. Note that as all elements of A are positive, implied by $\eta(A) \leq \gamma$, there exists ξ_0 such that P is a TPM for all $1 \leq \xi < \xi_0$. Also, $\eta_i(A) = \eta_i(P)$ for $i = 2, \dots, m$ and so, any difference in $\eta(A)$ and $\eta(P)$ comes from the difference between $\eta_1(A)$ and $\eta_1(P)$. Next, note that

$$\eta_1(P) = \max_{ij} \left\{ \frac{p_{1i}}{p_{1j}} \right\} = \max_{ij} \left\{ \frac{a_{1i} + (1 - \xi)a_{2i}}{a_{1j} + (1 - \xi)a_{2j}} \right\}$$

is a continuous function of ξ , and for $\xi = 1$, $\eta_1(P) = \eta_1(A) < \gamma$. So, there exists $1 < \xi < \xi_0$ for which $\eta_1(P) \leq \gamma$ and consequently, $\eta(P) \leq \gamma$. Take such a value ξ_* and use that in the construction of P .

Finally, note that $P = CA$ and $A = C^{-1}P$, with

$$C = \left(\begin{array}{cc|c} 1 & 1 - \xi_* & 0 \\ 0 & \xi_* & 0 \\ \hline 0 & 0 & I \end{array} \right) \quad \text{and} \quad C^{-1} = \left(\begin{array}{cc|c} 1 & 1 - 1/\xi_* & 0 \\ 0 & 1/\xi_* & 0 \\ \hline 0 & 0 & I \end{array} \right).$$

Now, as $\xi_* > 1$, C^{-1} is a TPM and hence $P \succeq A$. Also, as C is nonsingular, $P = DA$ only with $D = C$. But, C is not a TPM, as $\xi_* > 1$, and hence A is not sufficient for P . In summary, $P \succ A$ and hence A is inadmissible. \square

Lemma 4.3. *Suppose $A \in \mathcal{C}_\gamma$, $\eta_i(A)$ equals 1 or γ for all i , and $\eta_i(A) = 1$ for some i . Then, there exists $P \in \mathcal{C}_\gamma$ such that $P \succeq A$ and P satisfies the condition C1.*

Proof. Note that A can have at most one constant row because $A \in \mathcal{C}_\gamma$ and thus cannot have proportional rows. For notational simplicity, suppose that $\eta_1(A) = 1$, i.e., the all values in row 1 of A are the same, say δ . Then, from likelihood perspective, the response d_1 does not give any information about π . Intuitively, we may eliminate the response d_1 and distribute its probability (proportionally) to other responses. Specifically, construct P , from A , by deleting the first row and multiplying all other elements by $(1 - \delta)^{-1}$. It can be seen easily that row parity of the retained rows remain the same and $A_{m \times k} = C_{m \times (m-1)} P_{(m-1) \times k}$, where all elements of the first row of C are δ and the remaining rows constitute $(1 - \delta)I_{m-1}$. Thus, P satisfies C1 and $P \succeq A$. \square

If P as constructed in Lemma 4.3 also satisfies C2, i.e., $P \in \mathcal{C}_\gamma^0$, then from Corollary 4.1, it follows that $P \succ A$ and hence A is inadmissible. As we shall show next, if P does not satisfy C2, then P is inadmissible, which implies A is also inadmissible. Note that together Lemmas 4.2 and 4.3 cover all forms of violations of C1. The following lemma also completes the proof of Theorem 4.3, stated below.

Lemma 4.4. *Any randomization procedure $A \in \mathcal{C}_\gamma$ that satisfies the condition C1 but not C2 is inadmissible within \mathcal{C}_γ .*

Proof. Suppose $A_{m \times k} \in \mathcal{C}_\gamma$ satisfies C1 but not C2. Thus, $\eta_i(A) = \gamma$ for $i = 1, \dots, m$, and at least one row of A contains more than two distinct values. For notational simplicity, suppose the first row contains three or more distinct values and a_{11} is a ‘‘middle’’ value, i.e., $t < a_{11} < T$, where $t = \min_i \{a_{1i}\}$

and $T = \max_i \{a_{1i}\}$. Note that $T/t = \gamma$ as A satisfies C1. Let $\delta = (T - a_{11})/(T - t)$. Consider $P_{(m+1) \times k}^*$ whose rows are: $\vec{p}_1 = \delta(t, a_{12}, \dots, a_{1k})$, $\vec{p}_2 = (1 - \delta)(T, a_{12}, \dots, a_{1k})$ and $\vec{p}_i = \vec{a}_{i-1}$, $i = 3, \dots, m + 1$. It can be verified easily that P^* is a TPM, $\eta_i(P^*) = \gamma$ for $i = 1, \dots, m + 1$ and $A = CP^*$, with $C = \left(\begin{array}{cc|c} 1 & 1 & 0 \\ 0 & 0 & T \end{array} \right)$ and thus $P^* \succeq A$. Repeat the process to eliminate all “middle” values of A and if it creates any proportional rows, add those as per Remark 4.1. The resulting P belongs to \mathcal{C}_γ^0 and $P \succeq A$. Finally, in view of Corollary 4.1, we can conclude that $P \succ A$ and thus A is inadmissible. \square

Theorem 4.3. *A randomization procedure $P \in \mathcal{C}_\gamma$ is admissible within \mathcal{C}_γ only if P satisfies C1 and C2, i.e., $P \in \mathcal{C}_\gamma^0$.*

4.2. Optimality results

Generally, the class \mathcal{C}_γ^0 of all admissible procedures contains many P . However, for $k = 2$, it can be seen easily that if $P_{m \times 2}$ satisfies the condition C1 and has no proportional rows, then we must have $m = 2$. Moreover, \mathcal{C}_γ^0 consists of only two TPMs, which are also equivalent (by permutation). Thus, both are optimal procedures, one of which is reported below.

Proposition 4.1. *For binary X (i.e., $k = 2$), an optimal procedure at privacy level γ is given by $m = 2$, $p_{11} = p_{22} = \gamma(\gamma + 1)^{-1}$ and $p_{12} = p_{21} = (\gamma + 1)^{-1}$.*

For $k \geq 3$, choosing an optimal procedure from \mathcal{C}_γ^0 requires specific utility (or loss) functions. For a wide class of utility functions, Kairouz et al. (2016b) showed that under ϵ -DLP (which is equivalent to $\eta(P) \leq e^\epsilon$), an optimal procedure, under given π , can be obtained by solving a linear programming problem. Kairouz et al. (2016a) proved a close version of our Proposition 4.1. Duchi et al. (2018) obtained bounds on minimax risks. In the following, we shall present one result in a common setting.

Frequently, the categories of the survey variable are used as possible response categories, i.e., $m = k$ and $d_i = c_i$, $i = 1, \dots, k$, and consequently $\mathcal{S}_Z = \mathcal{S}_X$. In such cases, a common desire is to retain the original values as much as possible while meeting the privacy requirement. One mathematical formulation of this idea is to choose $P_{k \times k}$ to maximize $\sum_i p_{ii}$, the trace of P , subject to $\eta(P) \leq \gamma$, where γ is specified. The optimal P for this objective is given below.

Theorem 4.4. *Suppose $P_{k \times k}$ is a TPM with $\eta(P) \leq \gamma$. Then,*

- (a) $\sum_{i=1}^k p_{ii} \leq \frac{\gamma^k}{\gamma + k - 1}$ and
- (b) P attains the upper bound in (a) if and only if $p_{ii} = \frac{\gamma}{\gamma + k - 1}$ for all i and $p_{ij} = \frac{1}{\gamma + k - 1}$ for all $i \neq j$.

Proof. Take any $P_{k \times k}$ satisfying $\eta(P) \leq \gamma$, which implies that $p_{ii} \leq \gamma p_{ij}$ for all $i \neq j$. For fixed i , summing over $j \neq i$ and then adding p_{ii} to both sides, we get

$$(\gamma - 1 + k)p_{ii} \leq \gamma \sum_{j=1}^k p_{ij} \quad \text{or} \quad p_{ii} \leq \frac{\gamma}{\gamma + k - 1} \sum_{j=1}^k p_{ij}.$$

Then, adding both sides of the last inequality over i , and using the fact that for each j , $\sum_{i=1}^k p_{ij} = 1$, we obtain the inequality in (a).

The “if” part of (b) is easy to verify. For the “only if” part, the chain of inequalities in the preceding proof shows that equality in (a) holds if and only if $p_{ij} = p_{ii}/\gamma$ for all $i \neq j$. This implies that $p_{ij} = a_i/\gamma$ for all $i \neq j$, where a_1, \dots, a_k denote the diagonal elements of P . Now, as each column of P adds to 1, i.e., $a_j + \frac{1}{\gamma} \sum_{i \neq j} a_i = 1$ we obtain:

$$a_j \left(1 - \frac{1}{\gamma}\right) + \frac{1}{\gamma} \sum_{i=1}^k a_i = 1 \quad \text{or} \quad a_j = \frac{\gamma}{1 - \gamma} \left[1 - \sum_{i=1}^k a_i\right]$$

for all $j = 1, \dots, k$. Thus, we must have $a_1 = \dots = a_k$ and hence $p_{ii} = \frac{\gamma}{\gamma+k-1}$ for all i and $p_{ij} = \frac{1}{\gamma+k-1}$ for all $i \neq j$, as each column of P must add to 1. \square

Let P_0 denote the optimal TPM (for given k and γ) given above. Thus, the elements of P_0 are: $p_{ii} = \frac{\gamma}{\gamma+k-1}$ for all i and $p_{ij} = \frac{1}{\gamma+k-1}$ for all $i \neq j$. This P_0 has some attractive features and has received much attention. Note that P_0 is in \mathcal{C}_γ^0 and hence admissible. Agrawal et al. (2009) refer to P_0 as “the Gamma-Diagonal matrix” due to its structure; it has a common diagonal value and also a common off-diagonal value. They also proved an optimality property of P_0 , in terms of lowest condition number, among all symmetric positive definite P with $\eta(P) \leq \gamma$. Kairouz et al. (2016b) refer to P_0 as “the randomized response mechanism” and present certain mutual information optimality of P_0 .

5. Discussion

In this paper, we investigated the logic underlying the ρ_1 -to- ρ_2 and β -factor privacy criteria in full generality. We gave new insight and clarity using geometrical representation of privacy breach regions. We introduced the concepts of precise PBR and canonical strict information privacy to accurately describe the privacy demands of any stated criterion. Our Theorem 3.1, which gives necessary and sufficient conditions for attaining desired privacy, is a significant result. It also yields a numerical measure of the privacy demand of any given PBR, and shows that the parity of an RR procedure determine its privacy guarantee. It also gives a set of practically relevant PBRs and tells us to choose one of those in setting privacy requirement in real applications.

We compared data utility of privacy satisfying RR procedures using sufficiency of experiments, which is a strong criterion that does not rely on any specific loss function or utility measure. The class of all privacy preserving admissible RR procedures, \mathcal{C}_γ^0 , is an important finding. We also obtained the optimum procedure under a specific criterion, viz., maximize the trace of P subject to privacy constraints.

We believe that the requirement of no privacy breach for any property Q is overly stringent. Cell collapsing (or generalization) is a common privacy protection tool, which can be viewed as a special case of RR, with $P_\alpha(Z = d_i|X =$

$c_j) = 1$ if c_j is collapsed within d_i (or d_i contains c_j) and 0 otherwise. But, the parity of any such TPM is infinity, unless $m = 1$, in which case data utility is null. So, cell collapsing cannot give any strict information privacy without totally destroying data utility. It will be useful to modify the criterion by requiring no privacy breach for a subset of properties \mathcal{Q} but for all priors. We leave choosing \mathcal{Q} and appropriately modifying our results as future research topics.

6. Appendix: proofs

Proof of Lemma 3.2. The ‘only if’ part of the lemma follows readily. So, we shall prove only the ‘if’ part. Suppose (3.4) holds. Now, take any α , $Q \subseteq \{c_1, \dots, c_k\}$ and d_i such that $P_\alpha(Q) > 0$ and $P_\alpha(Z = d_i) > 0$. Suppose $c_q \in Q$ is such that $p_{iq} \geq p_{ij}$ for all j such that $c_j \in Q$. Consider the prior $\tilde{\alpha}$ with elements: $\tilde{\alpha}_j = \alpha_j$ if $c_j \notin Q$, $\tilde{\alpha}_q = P_\alpha(Q)$, and $\tilde{\alpha}_j = 0$ for all other j . Then, we have $P_{\tilde{\alpha}}(X = c_q) = \tilde{\alpha}_q = P_\alpha(Q)$ and

$$\begin{aligned} P_\alpha(Q|Z = d_i) &= \frac{\sum_{j:c_j \in Q} \alpha_j p_{ij}}{\sum_{j:c_j \in Q} \alpha_j p_{ij} + \sum_{j:c_j \notin Q} \alpha_j p_{ij}} \\ &\leq \frac{p_{iq}(\sum_{j:c_j \in Q} \alpha_j)}{p_{iq}(\sum_{j:c_j \in Q} \alpha_j) + \sum_{j:c_j \notin Q} \alpha_j p_{ij}} \\ &= \frac{\tilde{\alpha}_q p_{iq}}{\sum_{j=1}^k \tilde{\alpha}_j p_{ij}} = P_{\tilde{\alpha}}(X = c_q|Z = d_i) \\ &\leq h(P_{\tilde{\alpha}}(X = c_q)) = h(P_\alpha(Q)), \end{aligned}$$

where the first inequality holds by the fact that for $a > 0$, $\psi(x) = \frac{x}{x+a}$ is an increasing function of x over $(0, \infty)$ and the second inequality follows from (3.4). \square

Proof of Lemma 3.3. We shall prove that P satisfies (3.4) if and only if $\eta_i(P) \leq B(h)$ for $i = 1, \dots, m$. Take any (fixed) i , and note that (3.4) holds if $\alpha_j = 0$ or 1, or $p_{ij} = 0$. Take any j such that $p_{ij} > 0$ (which exists as each row of P contains at least one nonzero element). Then, for $0 < \alpha_j < 1$, we can write:

$$P_\alpha(X = c_j|Z = d_i) = \frac{\alpha_j p_{ij}}{\sum_{l=1}^k \alpha_l p_{il}} = \left[1 + \left(\frac{1 - \alpha_j}{\alpha_j} \right) \frac{1}{p_{ij}} \sum_{l:l \neq j} \left(\frac{\alpha_l}{1 - \alpha_j} \right) p_{il} \right]^{-1}. \quad (6.1)$$

In view of (6.1), for our fixed i and chosen j , (3.4) holds for all α if and only if

$$\sum_{l:l \neq j} \left(\frac{\alpha_l}{1 - \alpha_j} \right) \left(\frac{p_{il}}{p_{ij}} \right) \geq \left(\frac{\alpha_j}{1 - \alpha_j} \right) \left(\frac{1 - h(\alpha_j)}{h(\alpha_j)} \right) \quad (6.2)$$

for all α such that $0 < \alpha_j < 1$.

Letting $w_l = \frac{\alpha_l}{1 - \alpha_j}$ for $l \neq j$, it is seen that (6.2) is equivalent to

$$\sum_{l:l \neq j} w_l \left(\frac{p_{il}}{p_{ij}} \right) \geq \left(\frac{\alpha_j}{1 - \alpha_j} \right) \left(\frac{1 - h(\alpha_j)}{h(\alpha_j)} \right)$$

for all $0 < \alpha_j < 1$ and all $\{w_l\}$ such that $0 \leq w_l \leq 1$ and $\sum_{l \neq j} w_l = 1$. This holds if and only if

$$\inf_{\{w_l\}} \left(\sum_{l:l \neq j} w_l \left(\frac{p_{il}}{p_{ij}} \right) \right) \geq \sup_{0 < p < 1} \left(\frac{p}{1-p} \right) \left(\frac{1-h(p)}{h(p)} \right). \tag{6.3}$$

The infimum in (6.3) is $\min\{\frac{p_{il}}{p_{ij}} \mid l = 1, \dots, k, l \neq j\}$. Moreover, it must be positive in order to satisfy (6.3), because $h(p)$ cannot be 1 for all $0 < p < 1$ and hence right side of (6.3) is positive. This implies that p_{ij} must be positive for all $j = 1, \dots, k$. So, for our fixed i , (3.4) holds for all j and α if and only if (6.3) holds for $j = 1, \dots, k$, or equivalently,

$$\min\left\{ \frac{p_{il}}{p_{ij}} \mid j, l = 1, \dots, k \right\} \geq \sup_{0 < p < 1} \left(\frac{p}{1-p} \right) \left(\frac{1-h(p)}{h(p)} \right). \tag{6.4}$$

Note that both sides of (6.4) are positive and finite, and the above inequality can be recognized as $[\eta_i(P)]^{-1} \geq [B(h)]^{-1}$, which yields $\eta_i(P) \leq B(h)$. \square

Proof of Theorem 4.1. The ‘if’ part follows easily as noted in Remark 4.2. To prove the ‘only if’ part, suppose A and P are equivalent, i.e., there exist two TPMs $C_{m \times r}$ and $C_{r \times m}^*$ such that $A = CP$ and $P = C^*A$. Then, $(C^*C)_{r \times r}$ and $(CC^*)_{m \times m}$ are TPMs. Also, $C^*CP = C^*A = P$ or $(C^*C - I)P = 0$, and similarly, $(CC^* - I)A = 0$. These imply, by Lemma 6.1 (given below), that both (CC^*) and (C^*C) are identity matrices and consequently we must have $m = r$ (since both of them are of full rank) and $C^{-1} = C^*$. Now, since both are TPMs, C must be a permutation matrix (Minc, 1988, p. 3). \square

Lemma 6.1. *Suppose $B_{m \times m}$ and $P_{m \times k}$ are two transition probability matrices, P has no zero or proportional rows and $(B - I)P = 0$. Then, $B = I$, the identity matrix of order m .*

We shall use the following concepts and results to prove this lemma.

Definition 6.1. (Chakravarti, 1975) *A square matrix $B_{m \times m}$ is said to be reducible if there exists a permutation matrix Q such that*

$$Q^{-1}BQ = \begin{bmatrix} R & 0 \\ L & N \end{bmatrix}, \tag{6.5}$$

where R and N are square matrices. Otherwise, B is called irreducible.

If B is the TPM of a Markov chain, then B is irreducible means that one can always find a path between any two states. Note that $Q^{-1}BQ$ permutes the diagonal entries of B by exchanging the corresponding row and columns. We call this *diagonal permutation* in the following. In (6.5), if R or N are still reducible, they can be further reduced to the above form through diagonal permutation. Actually, if B is reducible, then through diagonal permutation, we can get a block lower-triangular matrix with irreducible diagonal blocks.

Theorem 6.1. (Chakravarti, 1975) If a non-negative matrix $B_{m \times m} = ((b_{ij}))$ is irreducible, then the matrix $F = B - D(r)$ must have rank $m - 1$, where $D(r)$ is the diagonal matrix with entries (r_1, r_2, \dots, r_m) , and $r_j = \sum_{i=1}^m b_{ij}$.

Definition 6.2. (Taussky, 1949) The column j of a square matrix $B = ((b_{ij}))$ is called weakly diagonal dominant, if $\sum_{i \neq j} |b_{ij}| \leq |b_{jj}|$. It is called strictly diagonal dominant if ' $<$ ' holds.

Theorem 6.2. (Taussky, 1949) Suppose $B_{m \times m}$ is an irreducible matrix, and all columns of B are weakly diagonal dominant and at least one is strictly diagonal dominant. Then, B is nonsingular.

Proof of Lemma 6.1. First, suppose $B_{m \times m}$ is irreducible, if possible. Let V_0 denote the vector space that is orthogonal to the row space of $(B - I)$. Note that if $(B - I)P = 0$, then all columns of P must be in V_0 . Applying Theorem 6.1 to B , noting that each column of B adds to 1 as B is a TPM, we obtain $\text{rank}(B - I) = m - 1$. This implies that the dimension of V_0 is 1 and hence all columns of P are proportional. Actually, they are identical (as the sum of each column is 1) and hence all rows of P are also identical. This contradicts the assumption that P has no proportional rows. Thus, B cannot be irreducible.

Next, suppose B is reducible. Then, there exists a permutation matrix Q such that $Q^{-1}BQ$ is a block lower-triangular matrix with irreducible diagonal blocks R_1, R_2, \dots, R_g . If all of these blocks are 1×1 identity matrices, then $B = I$ as $Q^{-1}BQ$ is TPM. If not, suppose R_{t+1} with dimension $s \times s$ is the first block that is not 1×1 identity matrix. This implies all the off-diagonal entries on the first t columns of $Q^{-1}BQ$ must be 0 (when $t \geq 1$). Take such a Q and denote $R = R_{t+1} = ((r_{ij}))$ to obtain

$$Q^{-1}BQ = \begin{bmatrix} I_t & 0 & 0 \\ 0 & R_{s \times s} & 0 \\ 0 & L & N \end{bmatrix}. \quad (6.6)$$

Note that $(B - I)P = 0$ implies that $QQ^{-1}(B - I)QQ^{-1}P = 0$ or

$$(Q^{-1}BQ - I)P^* = 0, \quad (6.7)$$

where $P^* = Q^{-1}P$, which is also a TPM with no proportional rows. In view of (6.6), equation (6.7) implies that $(R - I)P_s = 0$, where P_s consists of the $t + 1$ to $t + s$ rows of P^* . Here, each column of P_s is orthogonal to the rows of $(R - I)$, and hence must be in V_1 , the orthogonal space to the row space of $(R - I)$. We shall consider two cases to examine $\text{rank}(R - I)$ and its implication.

(i) $L = 0$ or L does not exist (i.e., $s = m - t$). Here, R is a TPM. Also, $s \geq 2$, since R is not 1×1 identity matrix. Apply Theorem 6.1 to R and the arguments used earlier (for irreducible B) to see that $\text{rank}(R - I) = s - 1$. So, the dimension of V_1 is 1, implying that all columns of P_s are constant multiples of a common vector and consequently all rows of P_s are the same. This contradicts the fact that P^* has no proportional rows. So, L cannot be a null matrix.

(ii) $L \neq 0$. Here, we shall apply Theorem 6.2 to $R^* = (R - I) = ((r_{ij}^*))$. First, R^* is irreducible, as R is so. Next, as each column of the right side of

(6.6) adds to 1, we get $\sum_{i=1}^s r_{ij} \leq 1$ for $j = 1, \dots, s$, and ' $<$ ' holds for at least one j , as $L \neq 0$. This shows, in view of $0 \leq r_{ij} \leq 1$, $r_{ii}^* = r_{ii} - 1$ and $r_{ij}^* = r_{ij}$ for $i \neq j$, that $\sum_{i \neq j} |r_{ij}^*| \leq |r_{jj}^*|$, for $j = 1, \dots, s$ and ' $<$ ' holds for some j . Thus, R^* satisfies the conditions of Theorem 6.2 and hence $\text{rank}(R^*) = s$, i.e., $(R - I)$ is nonsingular. Now, $(R - I)P_s = 0$ implies that $P_s = 0$, which is a contradiction. From the above discussion of all possible cases we must conclude that $B = I$. \square

Acknowledgment

We thank Dr. Martin Klein for some helpful remarks on an earlier draft. We also thank a referee for some helpful suggestions.

References

- Aggarwal, C.C. and Yu, P.S. (Eds.) (2008). *Privacy-Preserving Data Mining: Models and Algorithms*, New York: Springer Science and Business Media. [MR3308974](#)
- Agrawal, S., Haritsa, J.R. and Prakash, B.A. (2009). FRAPP: A Framework for high-accuracy privacy-preserving mining. *Data Mining and Knowledge Discovery*, 18, 101-139. [MR2469593](#)
- Basu, D. (1988). Likelihood and partial likelihood. In *Statistical Information and Likelihood: A Collection of Critical Essays by Dr. D. Basu*, J.K. Ghosh (ed.), Springer, New York, pp. 313-320. [MR0953081](#)
- Blackwell, D. (1951). Comparison of experiments. In *Proceedings of Second Berkeley Symposium on Mathematical Statistics and Probability*. University of California Press, Berkeley, pp. 93-102. [MR0046002](#)
- Blackwell, D. (1953). Equivalent comparison of experiments. *Annals of Mathematical Statistics*. 24, 265-272. [MR0056251](#)
- Boreale, M., and Paolini, M. (2015). Worst-and average-case privacy breaches in randomization mechanisms. *Theoretical Computer Science*, 597, 40-61. [MR3375747](#)
- Chakravarti, I.M. (1975). On a characterization of irreducibility of a non-negative matrix. *Linear Algebra and Its Applications*, 10, 103-109. [MR0393080](#)
- Chaudhuri, A. (2010). *Randomized Response and Indirect Questioning Techniques in Surveys*. Boca Raton: CRC Press. [MR2759226](#)
- Chaudhuri, A. and Mukerjee, R. (1988). *Randomized Response: Theory and Techniques*. New York: Marcel Dekker. [MR0933555](#)
- Chen, B-C., Kifer, D., LeFevre, K. and Machanavajjhala, A. (2009) Privacy-preserving data publishing. *Foundations and Trends in Databases*, 2, 1-167.
- Cruyff, M.J., Van Den Hout, A., and Van Der Heijden, P.G. (2008). The analysis of randomized response sum score variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70, 21-30. [MR2412629](#)
- Duchi, J.C., Jordan, M.I., and Wainwright, M.J. (2018). Minimax optimal pro-

- cedures for locally private estimation. *Journal of the American Statistical Association*, 113, 182-201. [MR3803452](#)
- Erlingsson, U., Pihur, V. and Korolova, A. (2014). Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, Scottsdale, Arizona, pp. 1054-1067.
- Evfimievski, A., Gehrke, J. and Srikant, R. (2003). Limiting privacy breaches in privacy-preserving data mining. *Proceedings of the 22nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS)*, San Diego, pp. 211-222.
- Evfimievski, A., Srikant, R., Agrawal, R. and Gehrke, J. (2004) Privacy preserving mining of association rules. *Information Systems*, 29, 343-364.
- Fung, B.C.M., Wang, K., Chen, R. and Yu, P.S. (2010). Privacy-preserving data publishing: A survey of recent developments. *ACM Computing Surveys*, 42, 14.
- Gouweleew, J.M., Kooiman, P., Willenborg, L.C.R.J. and De Wolf, P.-P. (1998). Post randomisation for statistical disclosure control: Theory and implementation. *Journal of Official Statistics*, 14, 463-478.
- Hundepool, A., Domingo-Ferrer, J., Franconi, L., Giessing, S., Nordholt, E.S., Spicer, K. and de Wolf, P.-P. (2012). *Statistical Disclosure Control*. New York: John Wiley & Sons. [MR3026260](#)
- Kairouz, P., Bonawitz, K., and Ramage, D. (2016a). Discrete distribution estimation under local privacy. In *Proceedings of the 33rd International Conference on Machine Learning*, New York, pp. 2436-2444.
- Kairouz, P., Oh, S., and Viswanath, P. (2016b). Extremal Mechanisms for Local Differential Privacy. *Journal of Machine Learning Research*, 17, 1-51. [MR3491111](#)
- Kass, R.E., and Raftery, A.E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90, 773-795. [MR3363402](#)
- Kifer, D. and Lin, B-R. (2012). An axiomatic view of statistical privacy and utility. *Journal of Privacy and Confidentiality*, 4, 5-49.
- Minc, H. (1988). *Nonnegative Matrices*. New York: John Wiley & Sons. [MR0932967](#)
- Nayak, T.K. and Adeshiyan, S.A. (2016). On invariant post-randomization for statistical disclosure control. *International Statistical Review*, 84, 26-42. [MR3491277](#)
- Nayak, T.K., Adeshiyan, S.A. and Zhang, C. (2016). A Concise Theory of Randomized Response Techniques for Privacy and Confidentiality Protection. *Handbook of Statistics*, 34, 273-286.
- Nayak, T.K., Zhang, C., and Adeshiyan, S.A. (2015). Emerging applications of randomized response concepts and some related issues. *Model Assisted Statistics and Applications*, 10, 335-344.
- Nayak, T.K., Zhang, C., and You, J. (2018). Measuring Identification Risk in Microdata Release and Its Control by Post-randomisation. *International Statistical Review*, 86, 300-321. [MR3852413](#)
- Taussky, O. (1949). A recurring theorem on determinants. *The American Math-*

- ematical Monthly*, 56, 672-676. [MR0032557](#)
- Torra, V. (2017). *Data Privacy: Foundations, New Developments and the Big Data Challenge*. New York: Springer.
- Van den Hout, A., and Elamir, E.A. (2006). Statistical disclosure control using post randomisation: Variants and measures for disclosure risk. *Journal of Official Statistics*, 22, 711-731.
- Van den Hout, A. and Van der Heijden, P.G. (2002). Randomized response, statistical disclosure control and misclassification: A review. *International Statistical Review*, 70, 269-288. [MR3006148](#)
- Warner, S.L. (1965). Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60, 63-69.
- Willenborg, L.C.R.J. and De Waal, T. (2001). *Elements of Statistical Disclosure Control*. New York: Springer. [MR1866909](#)