Journal of Privacy and Confidentiality
Vol. 10 (2) 2020

Submitted July 2018
Published June 2020

ON THE PRIVACY AND UTILITY PROPERTIES OF TRIPLE MATRIX-MASKING.

A. ADAM DING, GUANHONG MIAO, AND SAMUEL S. WU

Department of Mathematics, Northeastern University, Boston, MA
e-mail address: a.ding@neu.edu

Department of Biostatistics, University of Florida, Gainesville, FL.

Department of Biostatistics, University of Florida, Gainesville, FL.

Abstract. Privacy protection is an important requirement in many statistical studies. A recently proposed data collection method, triple matrix-masking, retains exact summary statistics without exposing the raw data at any point in the process. In this paper, we provide theoretical formulation and proofs showing that a modified version of the procedure is strong collection obfuscating: no party in the data collection process is able to gain knowledge of the individual level data, even with some partially masked data information in addition to the publicly published data. This provides a theoretical foundation for the usage of such a procedure to collect masked data that allows exact statistical inference for linear models, while preserving a well-defined notion of privacy protection for each individual participant in the study. This paper fits into a line of work tackling the problem of how to create useful synthetic data without having a trustworthy data aggregator. We achieve this by splitting the trust between two parties, the "masking service provider" and the "data collector."

1. Introduction

In the digital age, vast amount of data becomes available for research. At the same time, there is increasing pressure to protect the privacy of study subjects when their data is used. For medical research, the Health Insurance Portability and Accountability Act of 1996 and subsequent rulings have imposed legal requirements for privacy protection on the collection and handling of health data. Among other things, basic privacy protection measures include the removal of all personal identifiers when releasing data for use. However, simply removing the personal identifier variables does not prevent possible identification of the individual from other variables. To prevent the identification of an individual record, researchers have shown that released data should be aggregated to satisfy privacy conditions such as k-anonymity [Sweeney, 2002], l-diversity [Machanavajjhala et al., 2007] and t-closeness [Li et al., 2007].

However, releasing data only at aggregated levels severely restricts its usefulness in many research studies. Alternatively, methods are designed to release obfuscated micro-data

Key words and phrases: Data masking, privacy protection, matrix masking.

JPrivacy Confidentiality

www.journalprivacyconfidentiality.org
DOI:10.29012/jpc.674

that allows for the usual statistical analysis while preserving the privacy at individual levels. Some examples of such obfuscated micro-data publishing are: noise addition [Brand, 2002], multiple imputation[Rubin, 1993, Drechsler and Reiter, 2010], information preserving statistical obfuscation [Burridge, 2003], random projection based perturbation [Liu et al., 2006], random orthogonal matrix masking [Ting et al., 2008]. Particularly, in the random orthogonal matrix masking scheme, a masked data set $\boldsymbol{AX}$ is published, where $\boldsymbol{X}$ denotes the data matrix of real responses and $\boldsymbol{A}$ is a random orthogonal matrix. The published data $\boldsymbol{AX}$ keeps the exact values for sufficient statistics of linear models, thus allowing exact statistical inference for many standard data analysis methods [Ting et al., 2008, Wu et al., 2017b] while protecting privacy by denying the user's direct access to the raw data $\boldsymbol{X}$. While the above methods all protect the privacy of individual entries through publishing only the random perturbed micro-data, the privacy protection can be lost when multiple micro-data sets are combined from multiple inquires to the same database. Differential privacy is proposed to quantify the effectiveness of privacy protection of the random noise addition/perturbation schemes [Dwork et al., 2006, Dwork, 2006, Dwork and Naor, 2008] against multiple inquires to the database. Then the noise level can be adjusted to achieve a quantified tradeoff between inference efficacy and privacy preservation (measured by the differential privacy metric).

Traditionally, there is a trustworthy data collector/manager that collects raw data and ensure privacy protection by releasing the data sets with random perturbations. Such procedures however do not protect against attacks where an unscrupulous party has unauthorized access to the raw data set $\boldsymbol{X}$ kept by these centers. Such security breaks are becoming more common as shown by the recent well-publicized incidences involving hacking against databases at major retailers, banks and credit bureaus [Huffington Post, 2011, Reuters., 2015, 2017].

This paper fits into a line of work tackling the problem of how to create useful synthetic data without having a trustworthy data aggregator, and provides a theoretical study of the triple matrix-masking (TM^2) procedure [Wu et al., 2017b] that does not assume such a trustworthy data collector/manager. The TM^2 procedure is a multi-party collection and masking system that aims to collect and publish the random orthogonal masked data set $\boldsymbol{AX}$. We prove that, assuming no collusion between parties, no party learns more than the orbit of the data matrix under the action of the orthogonal group. More specifically, given the view of a particular party, let $\mathcal{S}$ be the set of data matrices that could possibly have resulted in that view. We show that $\mathcal{S}$ contains the full orbit of the data matrix and that given any prior on the data matrix, the party's posterior is simply their prior restricted to $\mathcal{S}$. We call data collection procedures with such properties as *strong obfuscating* since any extra information beyond $\boldsymbol{AX}$ available to a party does not help in further identifying the individual level data.

In the differential privacy literature, the issue of untrustworthy data collector can be dealt with using *local differential privacy* procedures [Kasiviswanathan et al., 2011], where noises are added to the individual data before passing to the data collector. The resulting synthetic data from differential privacy procedures, however, does not preserve exact statistics hence require special inference procedures designed to achieve optimal statistical inference Duchi et al. [2017]. Our TM^2 procedure provides an alternative where the published masked data exactly preserve any statistics of the data that are preserved under the action of the orthogonal group. This provides an useful utility that exact statistical inference for linear models are preserved, thus standard linear statistical inference procedures can be applied directly on the resulting synthetic data from the TM^2 procedure. On the other hand, the TM^2 procedure is only for a one-shot collection of each individual's data. When the individual data providers are sampled in multiple independent collections by different data collectors, differential privacy procedures can measure and limit the privacy leakage for the composition of the multiple collections. The TM^2 procedure does not consider the privacy leakage for the composition of the multiple collections.

Section 2 describe the TM^2 procedure and two new modifications to make it strong obfuscating. The theoretical analysis is provided in Section 3. Section 4 provides a summary and more detailed discussions for the relationship of the TM^2 procedure to the differential privacy and multi party computation methods.

2. The Masked Data Collection Procedure TM^{2}nd Its Modification

The privacy-preserving data collection scheme TM^2 was proposed first in Wu et al. [2017b] and later expanded by Wu et al. [2017a]. We describe our modified basic version of the TM^2 method here:

Step 1. The data collectors plan the data collection, create the database structure, program the data collection system. They randomly generate a $p \times p$ random orthogonal matrix $\boldsymbol{B}$, which is distributed to the participants' data collection devices.

Step 2. Each participant's data x_1 (a vector of dimension p_1) is collected and merged with Gaussian noise x_2 (of dimension p_2) into a vector $x = (x_1, x_2)$ of dimension $p = p_1 + p_2$. Then x is right multiplied by $\boldsymbol{B}$ on the participant's device, and only the resulting masked data $x\boldsymbol{B}$ leaves the device and is sent to the masking service provider.

Step 3. The masking service provider generates another $n \times n$ random orthogonal matrix $\boldsymbol{A}_2$. After receiving data from all participants, it combines the individual data $x\boldsymbol{B}$ into a $n \times p$ matrix $\boldsymbol{X}\boldsymbol{B}$, left multiplies by $\boldsymbol{A}_2$ and sends the doubly masked data $\boldsymbol{A}_2\boldsymbol{X}\boldsymbol{B}$ to the data collectors.

Step 4. The data collectors multiply $\boldsymbol{A}_2\boldsymbol{X}\boldsymbol{B}$ by $\boldsymbol{B}^{-1}$ to get back $\boldsymbol{A}_2\boldsymbol{X}$ and take the first p_1 columns to get $\boldsymbol{A}_2\boldsymbol{X}_1$. Then the data collectors generate another $n \times n$ random orthogonal matrix $\boldsymbol{A}_1$, left multiply it to $\boldsymbol{A}_2\boldsymbol{X}_1$, and publish $\boldsymbol{A}\boldsymbol{X}_1$ (where $\boldsymbol{A} = \boldsymbol{A}_1\boldsymbol{A}_2$) which is accessible by all data users.

Detailed theoretical analysis of the privacy guarantee on the TM^2 method has been missing. This paper fills that gap by proving theoretically that this modified version of the TM^2 method is strong obfuscating. We prove the strong obfuscating guarantee by showing that: (A) the extra information that any party in the process owns will not allow the party to reduce the data domain (possible values of data) small enough to identify individual level data; and (B) there is no statistical information leakage beyond the domain restriction considered in (A).

Compared to the original TM^2 scheme in Wu et al. [2017b], we make two modifications on the TM^2 procedure. For the first modification, we add random Gaussian noise in Step 2. The data collector wants to collect p_1 variables on n individuals so that the real response matrix becomes $\boldsymbol{X}_1$ of dimensions $n \times p_1$. We ask each participant to generate p_2 pure Gaussian noise variables, on his/her device according to a fixed variance parameter σ^2. Hence, the full data matrix would be $\boldsymbol{X} = (\boldsymbol{X}_1, \boldsymbol{X}_2)$. For privacy protections proved in later sections, we require that $p_1 < n \leq p = p_1 + p_2$. In Step 2 of this modified procedure, Gaussian noise x_2 is mingled with real response x_1 to provide protection in addition to the random mask $\boldsymbol{B}$. In Step 4, after the collectors get back $\boldsymbol{A}_2\boldsymbol{X} = (\boldsymbol{A}_2\boldsymbol{X}_1, \boldsymbol{A}_2\boldsymbol{X}_2)$, they

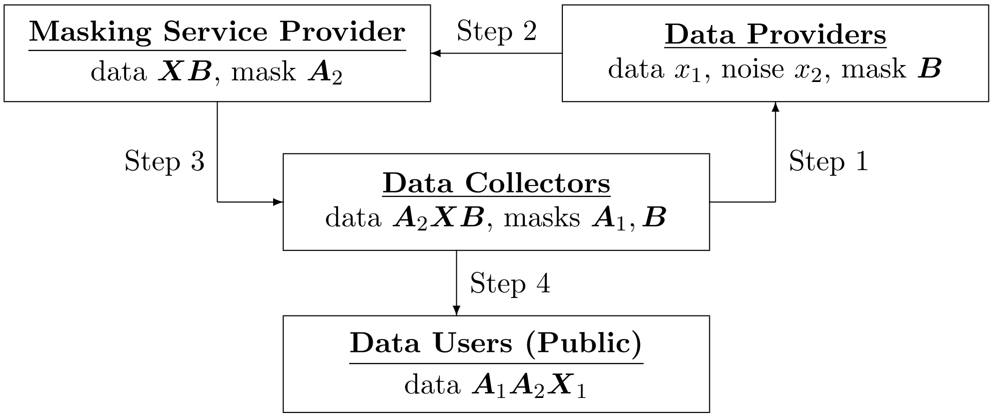
Masking Service Provider
data $\boldsymbol{XB}$, mask $\boldsymbol{A}_2$
Step 2
Data Providers
data x_1, noise x_2, mask $\boldsymbol{B}$
Step 3
Step 1
Data Collectors
data $\boldsymbol{A}_2\boldsymbol{XB}$, masks $\boldsymbol{A}_1, \boldsymbol{B}$
Step 4
Data Users (Public)
data $\boldsymbol{A}_1\boldsymbol{A}_2\boldsymbol{X}_1$

Figure 1: *This diagram shows each party's knowledge about the data and the masking matrices in the modified TM^2 method. Each party knows some masked version of the data:* $\boldsymbol{XB}$ *for the masking service provider,* $\boldsymbol{A}_2X$ *for the data collector, and* $\boldsymbol{A}_1\boldsymbol{A}_2\boldsymbol{X}_1$ *for everybody including the public. Nobody knows the original data* $\boldsymbol{X}_1$*, with each data provider (participant) knowing only his/her row* x_1

separate the matrix and discard those noises. Therefore the published data set $\boldsymbol{AX}_1$ with $\boldsymbol{A} = \boldsymbol{A}_1\boldsymbol{A}_2$ still gives the exact summary statistic, as it is only masked by $\boldsymbol{A}$ without containing the added noise.

For the second modification, instead of using a random invertible matrix for the right mask $\boldsymbol{B}$ as originally proposed by Wu et al. [2017b], we use a random orthogonal matrix for the right mask $\boldsymbol{B}$. As we will see in the privacy analysis in the next section, using an invertible matrix does make one part of the privacy proof easier. However, the other part of privacy proof depends on using a uniformly distributed random matrix to avoid information leakage that can lead to probabilistic attacks. While there is a natural uniform distribution on all orthogonal matrices that is well studied in literature, there is no natural uniform distribution on the set of all invertible matrices. The uniformly distributed orthogonal matrix $\boldsymbol{B}$ does provide sufficient privacy protection when combined with the addition of noise $\boldsymbol{X}_2$.

2.1. Privacy Analysis of TM^{2}etup.

To rigorously study the privacy protection issues in this data collection process, we analyze the information that can be accessed by each party and analyze whether such information allows inference of the individual level data.

First, we illustrate how to analyze the privacy protection assuming that the adversary only has access to the publicly published left-masked data set $\boldsymbol{M}_L = \boldsymbol{AX}_1$ where $\boldsymbol{A}$ is a random $n \times n$ orthogonal matrix. The issue becomes whether an adversary can identify individual level data knowing only $\boldsymbol{M}_L = \boldsymbol{m}_L$.

We consider the analysis in two stages. When given $\boldsymbol{M}_L = \boldsymbol{m}_L$, this restricts the possible values of $\boldsymbol{X}_1$ and can thus reveal information. In the first stage, we consider whether this support restriction on $\boldsymbol{X}_1$ (due to $\boldsymbol{M}_L = \boldsymbol{m}_L$) enables the identification of individual data. Let $\mathcal{S}_{\boldsymbol{X}_1}$ denote the support of $\boldsymbol{X}_1$, and let $\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{m}_L)$ denote the restricted support of $\boldsymbol{X}_1$ given that $\boldsymbol{M}_L = \boldsymbol{m}_L$. The privacy preservation depends on the size of $\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{m}_L)$. For example, in an extreme case, if the $\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{m}_L)$ contains only one matrix, then $\boldsymbol{X}_1$ is known to everyone and data privacy cannot be protected. Generally we can show, in next section, that this restricted support $\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{m}_L)$ is big enough so that identification of individual data is impossible.

In the second stage, we consider whether the adversary can learn any information beyond the restriction on support which was analyzed in the first stage. Such information can enable adversaries to launch probabilistic attacks [Machanavajjhala et al., 2007, Fung et al., 2010]. Fortunately, due to the independence between the mask $\boldsymbol{A}$ and the raw data $\boldsymbol{X}_1$, we can show that the posterior density of $\boldsymbol{X}_1$ given $\boldsymbol{M}_L = \boldsymbol{m}_L$ is the same as the prior density of $\boldsymbol{X}_1$ restricted to the support $\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{m}_L)$. Thus any loss of privacy is through the support restriction already studied in stage one. Therefore, knowing $\boldsymbol{M}_L = \boldsymbol{m}_L$ does not identify individual level data.

Next, we consider the privacy protection for all parties involved in the whole TM^2 data collection process. That is, we conduct the above two-stage privacy protection analysis given all information available to one party in the process. The data collector and the masking service provider each have access to some intermediate masked data in addition to the public data. Hence, we need to analyze privacy protection for an adversary knowing this intermediate masked data together with the public data set $\boldsymbol{M}_L = \boldsymbol{A}\boldsymbol{X}_1$.

The data collector knows, in addition to $\boldsymbol{M}_L = \boldsymbol{A}\boldsymbol{X}_1$, the double masked data $\boldsymbol{A}_2\boldsymbol{X}\boldsymbol{B}$. Since the data collector knows the masks $\boldsymbol{A}_1$ and $\boldsymbol{B}$, knowing these data $\boldsymbol{A}_1\boldsymbol{A}_2\boldsymbol{X}_1$ and $\boldsymbol{A}_2\boldsymbol{X}\boldsymbol{B}$ are simply equivalent to knowing $\boldsymbol{A}_2\boldsymbol{X}$. Due to the fact that $\boldsymbol{X}_2$ is purely noise independent of raw data $\boldsymbol{X}_1$, the theoretical privacy analysis for the data collector knowing $\boldsymbol{A}_2\boldsymbol{X} = (\boldsymbol{A}_2\boldsymbol{X}_1, \boldsymbol{A}_2\boldsymbol{X}_2)$ will have basically the same results as the analysis for the user with access only to $\boldsymbol{M}_L = \boldsymbol{A}\boldsymbol{X}_1$.

The masking service provider has access to the right-masked data $\boldsymbol{M}_R = \boldsymbol{X}\boldsymbol{B}$ in addition to the public left-masked data $\boldsymbol{M}_L = \boldsymbol{A}\boldsymbol{X}_1$. This information results in the most severe restriction on the support when compared to what resulted from knowledge by other parties. Thus, this is the weakest link for privacy preservation in the whole TM^2 data collection scheme. In Section 3, we present details of the two-stage privacy protection analysis when both $\boldsymbol{M}_L$ and $\boldsymbol{M}_R$ are known.

3. Theoretical Analysis of Privacy Preservation of TM^2

3.1. Notations, Formalizations and Technical Preliminaries.

We denote the probability densities of random matrices $\boldsymbol{X}_1$, $\boldsymbol{X}_2$, $\boldsymbol{A}$ and $\boldsymbol{B}$ as $\pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1)$, $\pi_{\boldsymbol{X}_2}(\boldsymbol{x}_2)$, $\pi_{\boldsymbol{A}}(\boldsymbol{a})$ and $\pi_{\boldsymbol{B}}(\boldsymbol{b})$ respectively. The supports of these distributions are denoted respectively as $\mathcal{S}_{\boldsymbol{X}_1}$, $\mathcal{S}_{\boldsymbol{X}_2}$, $\mathcal{S}_{\boldsymbol{A}}$ and $\mathcal{S}_{\boldsymbol{B}}$.

We want to study, based on information $INFO$ available to one party, what this party can infer about the individual level data. Here this $INFO$ includes the publicly available final left-masked data $\boldsymbol{M}_L = \boldsymbol{A}\boldsymbol{X}_1$ and some extra information available to the particular party. The restricted support of $\boldsymbol{X}_1$ given $INFO$ is denoted as $\mathcal{S}_{\boldsymbol{X}_1}(INFO)$ which consists of all $n \times p$ matrices that can be the value of $\boldsymbol{X}_1$ which is compatible with $INFO$.

For example, given only the public masked data $INFO = \boldsymbol{M}_L$, the restricted support is

$$\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{M}_L) = \{\boldsymbol{U} : \exists \tilde{\boldsymbol{A}} \in \mathcal{S}_{\boldsymbol{A}} \text{ such that } \tilde{\boldsymbol{A}}\boldsymbol{U} = \boldsymbol{M}_L.\}$$

Let $\mathcal{O}_n$ denote the set consisting of all orthogonal matrices. In the case of left masking with a random orthogonal matrix $\boldsymbol{A}$, for any orthogonal $\bar{\boldsymbol{A}} \in \mathcal{O}_n$, $\boldsymbol{U} = \bar{\boldsymbol{A}}\boldsymbol{X}_1$ is compatible with $INFO = \boldsymbol{M}_L$. That is, $\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{M}_L) = \mathcal{O}_n\boldsymbol{X}_1$. To see this, let $\tilde{\boldsymbol{A}} = \boldsymbol{A}\bar{\boldsymbol{A}}^T$, then $\tilde{\boldsymbol{A}} \in \mathcal{O}_n$ and $\tilde{\boldsymbol{A}}(\bar{\boldsymbol{A}}\boldsymbol{X}_1) = \boldsymbol{M}_L$. Here and throughout this paper, we use ^{T} to denote the transpose of a matrix.

For the strong obfuscating guarantee, we wish to show that the extra information available to the parties in the process does not cause any privacy loss more than the publicly released final left-masked data $\boldsymbol{M}_L = \boldsymbol{A}\boldsymbol{X}_1$. We want to show that: stage one (i) the restricted support $\mathcal{S}_{\boldsymbol{X}_1}(INFO)$ is the same as $\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{M}_L) = \mathcal{O}_n\boldsymbol{X}_1$; stage two (ii) the conditional probability distribution of $\boldsymbol{X}_1$ given $INFO$ is similar to the probability distribution of $\boldsymbol{X}_1$ given support $\mathcal{S}_{\boldsymbol{X}_1}(INFO)$, thus there is no privacy loss through probability attacks beyond the loss from the support restriction considered in stage one.

We now formalize the precise mathematical statements to prove in stages one and two. More precisely for stage one, we hope that the restricted support is the same as if only the public left-masked data is available.

$$(i) \qquad \mathcal{S}_{\boldsymbol{X}_1}(INFO) = \mathcal{O}_n\boldsymbol{X}_1,$$

for $INFO$ available to any one party in the process. For the second stage, we denote $\pi_{\boldsymbol{X}_1|INFO}(\boldsymbol{x}_1|INFO)$ as the posterior distribution of $\boldsymbol{X}_1$ given $INFO$. The prior density $\pi_{\boldsymbol{X}_1}$ restricted on the support $\mathcal{S}_{\boldsymbol{X}_1}(INFO)$ is

$$\pi_{\boldsymbol{X}_1|\mathcal{S}_{\boldsymbol{X}_1}(INFO)}(\boldsymbol{x}_1) = \frac{\pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1)}{\int\limits_{\mathcal{S}_{\boldsymbol{X}_1}(INFO)} \pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1^*)d\boldsymbol{x}_1^*}.$$

To show that there is no extra privacy loss beyond the support restriction considered in stage one, we prove that these two probability densities agree with each other. That is, we wish to prove

$$(ii) \qquad \pi_{\boldsymbol{X}_1|INFO}(\boldsymbol{x}_1|INFO) = \pi_{\boldsymbol{X}_1|\mathcal{S}_{\boldsymbol{X}_1}(INFO)}(\boldsymbol{x}_1).$$

Definition 3.1. A data collection process is *strong collection obfuscating* if conditions (i) and (ii) hold for the information $INFO$ available to any party in this process.

A slightly weaker version is that the above property holds with a high probability. Notice that the $INFO$ available to any party in this process can be determined from the values of $\boldsymbol{X}_1$, $\boldsymbol{X}_2$, $\boldsymbol{A}$ and $\boldsymbol{B}$ which are generated respectively from distributions with densities $\pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1)$, $\pi_{\boldsymbol{X}_2}(\boldsymbol{x}_2)$, $\pi_{\boldsymbol{A}}(\boldsymbol{a})$ and $\pi_{\boldsymbol{B}}(\boldsymbol{b})$. Thus such $INFO$ is generated from a probability distribution defined by $\pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1)$, $\pi_{\boldsymbol{X}_2}(\boldsymbol{x}_2)$, $\pi_{\boldsymbol{A}}(\boldsymbol{a})$ and $\pi_{\boldsymbol{B}}(\boldsymbol{b})$. We want that, with high probability from this probability distribution, the generated values of $INFO$ satisfy conditions (i) and (ii).

Definition 3.2. A data collection process is *ϵ-strong collection obfuscating* if, with probability at least $1-\epsilon$, conditions (i) and (ii) hold for the information $INFO$ available to any party in this process.

Our definition of the strong collection obfuscating procedure ensures that there is no privacy loss due to observations by any party in the process beyond those contained in the publicly released final data. This definition delineates the privacy protection in the collection process from the privacy protection in the publicly releasing of final data $\boldsymbol{M}_L = \boldsymbol{A}\boldsymbol{X}_1$. The theoretical analysis concentrates on the soundness of the collection process.

Given the public left-masked data $\boldsymbol{M}_L = \boldsymbol{A}\boldsymbol{X}_1$, the statistics $\boldsymbol{X}_1^T\boldsymbol{X}_1$ are released to the user. The user has the first two exact statistical moments and statistical models, such as linear regression, can be fitted exactly as if the user has the raw data set $\boldsymbol{X}_1$. And the residuals are known up to an orthogonal matrix multiplication, therefore the usual statistical model diagnostics methods can also be carried out as if done on the raw data set.

For continuous data, the user cannot recover the individual level data since the user only sees a linear combination of all individuals' data, and there is no utilizable statistical distributional information other than the prior (population) density $\pi_{\boldsymbol{X}_1}$. This ensures the privacy of individual data.

In practice, the types of elements in $\boldsymbol{X}_1$ may also be known to the user. This can further restrict the support. We assume that the elements in data matrix $\boldsymbol{X}_1$ are all encoded as numerical values (e.g., "yes/no" answer to a question may be encoded as 1 and 0). We consider that the type of data in each column, either as continuous/discrete/binary, is public knowledge. Let $\mathcal{S}_j$ denote the support of the type of data in the j-th column of $\boldsymbol{X}_1$. For example, if data are continuous, then $\mathcal{S}_j = \mathcal{R}$; if the data are binary, then $\mathcal{S}_j = \{0, 1\}$; if the data are positive integers, then $\mathcal{S}_j = \{1, 2, ...\} = \mathcal{N}^+$. Knowing the type of data in each column would restrict the support of $\boldsymbol{X}_1$ to

$$\tilde{\mathcal{S}}_{\boldsymbol{X}_1} = \{\boldsymbol{U} : (\text{all entries of } j\text{-th column in } \boldsymbol{U}) \in \mathcal{S}_j \qquad j = 1, ..., p_1.\}.$$

Then with knowledge of both $INFO$ and types of data, the restricted support becomes the intersection of $\tilde{\mathcal{S}}_{\boldsymbol{X}_1}$ and $\mathcal{S}_{\boldsymbol{X}_1}(INFO)$,

$$\mathcal{S}_{\boldsymbol{X}_1}(INFO; TYPE) = \tilde{\mathcal{S}}_{\boldsymbol{X}_1} \cap \mathcal{S}_{\boldsymbol{X}_1}(INFO).$$

Let $\mathcal{P}_n$ denote the set of all permutation matrix $\boldsymbol{P}$. Since all permutation matrices are orthogonal and permutation does not change the type of elements, we have the following Lemma:

Lemma 3.3. *For any strong obfuscating data collection process,*

$$\mathcal{P}_n\boldsymbol{X}_1 \subseteq [\tilde{\mathcal{S}}_{\boldsymbol{X}_1} \cap \mathcal{O}_n\boldsymbol{X}_1] = \mathcal{S}_{\boldsymbol{X}_1}(INFO; TYPE).$$

Lemma 3.3 indicates that a strong collection obfuscating data collection process offers some privacy protection even when the data types are known. Since all permutations are in the $\mathcal{S}_{\boldsymbol{X}_1}(INFO; TYPE)$, any individual cannot be identified here without extra side information. It is not clear whether the type can be combined with some side information (such as data that a particular individual is a smoker) to reveal other individual level data. However, notice that any weakness in this aspect is inherently due to releasing the public data $\boldsymbol{A}\boldsymbol{X}_1$. Our strong collection obfuscating procedure ensures that no extra privacy loss is added during the process beyond the privacy loss in releasing $\boldsymbol{A}\boldsymbol{X}_1$.

As we discussed in the previous section, the party with the most information during the TM^2 process is the masking service provider who knows $INFO = (\boldsymbol{M}_L, \boldsymbol{M}_R)$. Therefore, in the next section, we study when (i) and (ii) hold for $INFO = (\boldsymbol{M}_L, \boldsymbol{M}_R)$. Here we first state some technical preliminary results on the characterization of the uniform distribution on orthogonal matrices. These preliminary results are used in studying the second stage condition (ii) later.

Under the matrix multiplication, the orthogonal matrices form a compact Hausdorff topological group $\mathcal{O}_n$. Therefore, there is a unique Haar measure $\mu(\cdot)$ on $\mathcal{O}_n$ such that the measure of the whole sample space $\mathcal{O}_n$ equals one. Then this Haar measure induces a natural uniform distribution on $\mathcal{O}_n$. See Chapter 2 of Zhang [2014] for a detailed technical equivalent characterization of the uniform distribution on $\mathcal{O}_n$. Since a Haar measure $\mu(\cdot)$ is invariant under the matrix multiplication, the uniform distribution is also invariant under the matrix multiplication.

Lemma 3.4. *Let $\pi_0(\cdot)$ denote the probability density function of the uniform distribution on $\mathcal{O}_n$. Then for any orthogonal matrix $\boldsymbol{A}_0 \in \mathcal{O}_n$,*

$$\pi_0(\boldsymbol{a}) = \pi_0(\boldsymbol{A}_0\boldsymbol{a}) = \pi_0(\boldsymbol{a}\boldsymbol{A}_0), \qquad \textit{for all } \boldsymbol{a} \in \mathcal{O}_n. \tag{3.1}$$

Also, the product of two uniformly distributed orthogonal matrices is also uniformly distributed.

Lemma 3.5. *If $\boldsymbol{A}_1 \sim \pi_0$ and $\boldsymbol{A}_2 \ \sim \pi_0$ are independent of each other, then their product $\boldsymbol{A} = \boldsymbol{A}_1\boldsymbol{A}_2$ also follows the uniform distribution π_0 on $\mathcal{O}_n$.*

The proof is straightforward and can be found in Chapter 2 of Zhang [2014].

In the TM^2 scheme, when the masking service provider and the data collector each generate a random orthogonal matrix $\boldsymbol{A}_1$ and $\boldsymbol{A}_2$ respectively according to π_0, then the mask $\boldsymbol{A} = \boldsymbol{A}_1\boldsymbol{A}_2$ for the publicly released data set is also uniformly distributed. In practice, the uniformly distributed random orthogonal matrices can be generated using algorithms described in Heiberger [1978], Anderson et al. [1987], Wu et al. [2017b].

3.2. **Restricted Support Given Knowledge of Masked Data Sets.** We first prove that condition (i) holds for $INFO = (\boldsymbol{M}_L, \boldsymbol{M}_R)$ when invertible matrices are used for right mask $\boldsymbol{B}$ as originally proposed by Wu et al. [2017b]. That is, $\boldsymbol{B} \in \mathcal{I}_n$, where $\mathcal{I}_n$ denote the set of all $n \times n$ invertible matrices. Condition (i) then becomes that all orthogonal transformations of $\boldsymbol{X}_1$ are contained in the restricted support

$$\begin{aligned}\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{M}_L, \boldsymbol{M}_R) = \{\boldsymbol{U} :\ & \exists \tilde{\boldsymbol{A}} \in \mathcal{S}_{\boldsymbol{A}}, \tilde{\boldsymbol{B}} \in \mathcal{S}_{\boldsymbol{B}} \text{ and } \tilde{\boldsymbol{U}} \in \mathcal{S}_{\boldsymbol{X}_2} \text{ such that} \\ & \tilde{\boldsymbol{A}}\boldsymbol{U} = \boldsymbol{M}_L \text{ and } (\boldsymbol{U}, \tilde{\boldsymbol{U}})\tilde{\boldsymbol{B}} = \boldsymbol{M}_R\}.\end{aligned} \tag{3.2}$$

Theorem 3.6. *Suppose $\mathcal{S}_{\boldsymbol{A}} = \mathcal{O}_n$ and $\mathcal{S}_{\boldsymbol{B}} = \mathcal{I}_p$, $p_1 \leq n \leq p$ and $\boldsymbol{X}$ is full rank (i.e., rank$(\boldsymbol{X}) = n$), then for any $\boldsymbol{P} \in \mathcal{O}_n$, $\boldsymbol{P}\boldsymbol{X}_1 \in \mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{M}_L, \boldsymbol{M}_R)$. In other words, $\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{M}_L, \boldsymbol{M}_R) = \mathcal{O}_n\boldsymbol{X}_1 = \mathcal{O}_n\boldsymbol{M}_L$.*

Proof. We need to show that, for any $\boldsymbol{P} \in \mathcal{O}_n$, $\boldsymbol{U} = \boldsymbol{P}\boldsymbol{X}_1 \in \mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{M}_L, \boldsymbol{M}_R)$.

Since $\boldsymbol{P} \in \mathcal{O}_n$ and $\boldsymbol{A} \in \mathcal{O}_n$, then $\tilde{\boldsymbol{A}} = \boldsymbol{A}\boldsymbol{P}^T \in \mathcal{O}_n$. Then we have

$$\tilde{\boldsymbol{A}}\boldsymbol{U} = \boldsymbol{A}\boldsymbol{P}^T\boldsymbol{P}\boldsymbol{X}_1 = \boldsymbol{A}\boldsymbol{X}_1 = \boldsymbol{M}_L. \tag{3.3}$$

Since $\boldsymbol{X}$ is full-rank and $n \leq p$, there exists a $(p-n) \times p$ matrix $\boldsymbol{X}^*$ such that $\begin{pmatrix} \boldsymbol{X} \\ \boldsymbol{X}^* \end{pmatrix}$ is full-rank and thus invertible. Since $\boldsymbol{P} \in \mathcal{O}_n$, $\begin{pmatrix} \boldsymbol{P}^T\boldsymbol{X} \\ \boldsymbol{X}^* \end{pmatrix}$ is also full-rank and invertible. Hence we can define an invertible matrix

$$\tilde{\boldsymbol{B}} = \begin{pmatrix} \boldsymbol{X} \\ \boldsymbol{X}^* \end{pmatrix}^{-1} \begin{pmatrix} \boldsymbol{P}^T\boldsymbol{X} \\ \boldsymbol{X}^* \end{pmatrix} \boldsymbol{B}.$$

Also let $\tilde{\boldsymbol{U}} = \boldsymbol{P}\boldsymbol{X}_2$ thus $(\boldsymbol{U}, \tilde{\boldsymbol{U}}) = \boldsymbol{P}\boldsymbol{X}$, we have

$$\begin{pmatrix} (\boldsymbol{U}, \tilde{\boldsymbol{U}}) \\ \boldsymbol{X}^* \end{pmatrix} \tilde{\boldsymbol{B}} = \begin{pmatrix} \boldsymbol{P} & 0 \\ 0 & \boldsymbol{I} \end{pmatrix} \begin{pmatrix} \boldsymbol{X} \\ \boldsymbol{X}^* \end{pmatrix} \tilde{\boldsymbol{B}} = \begin{pmatrix} \boldsymbol{P} & 0 \\ 0 & \boldsymbol{I} \end{pmatrix} \begin{pmatrix} \boldsymbol{P}^T\boldsymbol{X} \\ \boldsymbol{X}^* \end{pmatrix} \boldsymbol{B} = \begin{pmatrix} \boldsymbol{X} \\ \boldsymbol{X}^* \end{pmatrix} \boldsymbol{B},$$

where $\boldsymbol{I}$ is the identity matrix of size $(p-n) \times (p-n)$. The first p rows in the last equation are

$$(\boldsymbol{U}, \tilde{\boldsymbol{U}})\tilde{\boldsymbol{B}} = \boldsymbol{X}\boldsymbol{B} = \boldsymbol{M}_R. \tag{3.4}$$

(3.3) and (3.4) together means that $\boldsymbol{U}$ satisfies (3.2), thus $\boldsymbol{U}$ belongs to $\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{M}_L, \boldsymbol{M}_R)$. □

Theorem 3.6 states that condition (i) is satisfied when the $\boldsymbol{X}$ is full rank. In the original TM^2 scheme proposal, the full rank condition may or may not be satisfied because it is determined by the underlying probability distribution of $\boldsymbol{X}_1$ which is outside the control of the designer of this procedure. With the modification of extra noise matrix $\boldsymbol{X}_2$, we can ensure the full rank condition by specifying the noise generation mechanism. Particularly, we specify that each individual data provider generates a p_2-dimension noise vector with i.i.d. elements from a Gaussian distribution with $p_2 \geq n$. This will ensure with probability one that $\boldsymbol{X}$ is indeed full rank.

Remark 1. (Size of the right mask) *For privacy preservation, the size of right mask p has to be bigger than the data size n as assumed in Theorem 3.6. When $p < n$, some rows of $\boldsymbol{M}_R$ are linear dependent which provides further restriction on the support. We provide such a counter example in Appendix A to illustrate that such a restriction together with knowledge of data type can reveal individual level data.*

Above we considered the support restriction under the original TM^2 scheme proposal [Wu et al., 2017b] of invertible right mask, $\mathcal{S}_{\boldsymbol{B}} = \mathcal{I}_p$. However, unlike $\mathcal{O}_p$, $\mathcal{I}_p$ does not form a compact Hausdorff topological group. Therefore, there exists no uniform distribution on $\mathcal{I}_p$. Due to the non-uniformity of $\boldsymbol{B}$, the posterior distribution of $\boldsymbol{X}_1$ given $(\boldsymbol{M}_L, \boldsymbol{M}_R)$ leaks information beyond the support restriction, thus the second stage condition (ii) no longer holds. This makes the usage of random invertible right masks in the TM^2 scheme very tricky. It is unclear what distribution on $\mathcal{I}_p$ should be used to generate the random invertible $\boldsymbol{B}$.

Here, we consider the modification of the TM^2 scheme where the right mask $\boldsymbol{B}$ is a random orthogonal matrix generated from the uniform distribution π_0 on $\mathcal{O}_p$. We show that if the random noise $\boldsymbol{X}_2$ is large enough, then condition (i) still holds when the orthogonal right mask $\boldsymbol{B}$ is used.

Let $\lambda_{min}(\boldsymbol{M})$ and $\lambda_{max}(\boldsymbol{M})$ denotes the minimum and the maximum eigenvalues of a semi-positive definite matrix $\boldsymbol{M}$. The restricted support will remain big if the noise is large enough:

$$\lambda_{min}(\boldsymbol{M}_R\boldsymbol{M}_R^T - \boldsymbol{X}_1\boldsymbol{X}_1^T) = \lambda_{min}(\boldsymbol{X}_2\boldsymbol{X}_2^T) > \lambda_{max}(\boldsymbol{X}_1\boldsymbol{X}_1^T). \tag{3.5}$$

Now we have a result similar to Theorem 3.6.

Theorem 3.7. *Suppose $\mathcal{S}_{\boldsymbol{A}} = \mathcal{O}_n$ and $\mathcal{S}_{\boldsymbol{B}} = \mathcal{O}_p$, $p_1 \leq n \leq p$. If condition(3.5) holds, then $\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{M}_L, \boldsymbol{M}_R) = \mathcal{O}_n\boldsymbol{X}_1$.*

The proof is provided in Appendix B.

Next, we show that condition (ii) also holds when under condition(3.5). Then we discuss how achievable the technical condition(3.5) is in practice.

3.3. **Information Leakage Beyond the Support Restriction.** We now study the second stage condition (ii) by checking the amount of information an adversary can get from the posterior distribution of $\boldsymbol{X}_1$ given $(\boldsymbol{M}_L, \boldsymbol{M}_R)$ beyond their restriction on the support of $\boldsymbol{X}_1$. Given $INFO = (\boldsymbol{M}_L, \boldsymbol{M}_R)$, the posterior density is denoted as $\pi_{\boldsymbol{X}_1|(\boldsymbol{M}_L,\boldsymbol{M}_R)}(\boldsymbol{x}_1|\boldsymbol{m}_L, \boldsymbol{m}_R)$. The prior density $\pi_{\boldsymbol{X}_1}$ restricted on the support $\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{m}_L, \boldsymbol{m}_R)$ is denoted as $\pi_{\boldsymbol{X}_1|\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{m}_L,\boldsymbol{m}_R)}$.

Theorem 3.8. *Let $\boldsymbol{X}_1$ be a random matrix with probability density $\pi_{\boldsymbol{X}_1}$. We assume that the elements in $\boldsymbol{X}_2$ are generated i.i.d. from a Gaussian distribution with mean zero. When condition (3.5) holds, given $\boldsymbol{M}_L$ and $\boldsymbol{M}_R$, the posterior density of $\boldsymbol{X}_1$ is the same as the prior density restricted on $\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{M}_L, \boldsymbol{M}_R)$. That is,*

$$\pi_{\boldsymbol{X}_1|(\boldsymbol{M}_L,\boldsymbol{M}_R)}(\boldsymbol{x}_1|\boldsymbol{m}_L, \boldsymbol{m}_R) = \pi_{\boldsymbol{X}_1|\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{m}_L,\boldsymbol{m}_R)}(\boldsymbol{x}_1). \tag{3.6}$$

The proof of Theorem 3.8 is provided in Appendix D.

3.4. **ϵ-strong obfuscating TM^2.** Theorem 3.7 and Theorem 3.8 states, respectively, that conditions (i) and (ii) hold under condition (3.5). Combining them, we have the following Theorem.

Theorem 3.9. *If*

$$Pr[\lambda_{min}(\boldsymbol{X}_2\boldsymbol{X}_2^T) > \lambda_{max}(\boldsymbol{X}_1\boldsymbol{X}_1^T)] \geq 1 - \epsilon, \tag{3.7}$$

then the proposed TM^2 procedure is ϵ-strong collection obfuscating by Definition 3.2.

The ϵ-strong collection privacy property ensures that there is at most ϵ probability for the process to leak any privacy information beyond the public released data $\boldsymbol{A}\boldsymbol{X}_1$. TM^2 achieves this property when the technical condition (3.7) holds. To achieve the technical condition (3.7), we generate the p_2-dimensional noise vector x_2 with i.i.d. Gaussian elements of mean zero and a sufficiently large variance σ^2. We present a technical probability bound in Appendix C, where the probability of violating condition (3.5) is decreasing exponentially and specifics a σ^2 value which ensures condition (3.7). Larger variance σ^2 always increases the probability that condition (3.5) holds. In practice, the variance σ^2 is only limited due to the computation accuracy. That is, σ should not exceed raw data values by the orders of magnitude allowed by the machine precision.

3.5. **Extension to Alleviate Collusion Risks.** We have shown that the privacy of individual data can be protected when no party in the TM^2 scheme knows all the masks. However, there are also risks of collusion among different parties in the procedure. Since the right mask $\boldsymbol{B}$ is known to the data collector and all individual data providers, if one of them share this info with the masking services provider, then the privacy protection can be broken.

Wu et al. [2017a] proposed ideas to protect against this collusion risk using the ideas of multiparty computation. For each individual, the data vector x can be broken up as K_1 random components x^1,...,x^{K_1} where $x = x^1 + ... + x^{K_1}$. Then such components are sent to K_1 right masking service providers, one to each. The resulting masked data $x^i\boldsymbol{B}_i$, $i = 1, ..., K_1$, are then sent to the left masking service provider to be merged and created the double masked data $\boldsymbol{A}\boldsymbol{X}^i\boldsymbol{B}_i$, $i = 1, ..., K_1$. For further protection, they can be passed through K_2 left masking service providers to generate $\boldsymbol{A}_{K_2}...\boldsymbol{A}_1\boldsymbol{X}^i\boldsymbol{B}_i$. Let $\boldsymbol{A} = \boldsymbol{A}_{K_2}\boldsymbol{A}_{K_2-1}...\boldsymbol{A}_1$. Then the masked data $\boldsymbol{A}\boldsymbol{X}^i\boldsymbol{B}_i$, $i = 1, ..., K_1$, are sent to the corresponding right masking service providers to remove the right masking. Then the resulting $\boldsymbol{AX}^i$, $i = 1, ..., K_1$, are sent to the data collector to generate $\boldsymbol{AX} = \boldsymbol{AX}^1 + ... + \boldsymbol{AX}^{K_1}$. Unless all K_1 right (or all K_2 left) masking service providers collude, they cannot find values of all components $\boldsymbol{X}^1, ..., \boldsymbol{X}^{K_1}$.

The stage one theoretical analysis on this extended TM^2 scheme can be analyzed similarly as before, where the restricted support condition (i*) holds given condition (3.5). The stage two analysis is more involved, as the posterior distribution of $\boldsymbol{X}_1$ given some shares, depends on the distribution of the shares. Which random distributions should the shares be generated from to effect no additional privacy loss remains an open question, and will be investigated in future work.

4. Discussions and Conclusions

This paper conducts a theoretical analysis of privacy preservation in a modified TM^2 scheme. Random noises were used with uniformly distributed orthogonal matrix masks to hide individual data during the data collection process. The noise addition in the first step of the TM^2 scheme is similar to the idea of noise perturbed response schemes. However, the critical difference is that our noise addition is only intended to help mask data during the transition, and is in fact removed after the right mask removal. The resulting published data set is a left masked data set with exact summary statistics, unlike many other noise addition schemes where the summary statistics are randomly approximated.

This work is aimed to protect against unscrupulous access to the raw data $\boldsymbol{X}_1$ traditionally hold by a trusted operator. We would like to further clarify the relationship to differential privacy methods [Dwork et al., 2006] which aims to provide a strong privacy protection and closure under composition of multiple accesses to the database. There are two types of differential privacy models. In the central model, a trusted database operator holds the raw data, and releases noise perturbed summary statistics for inquires. In the local model [Evfimievski et al., 2003, Dwork et al., 2006, Kasiviswanathan et al., 2011, Cormode et al., 2018], noise is added at the individual level based on the idea of randomized response methods [Warner, 1965, Blair et al., 2015]. The local differential privacy procedures similarly addresses the issue of untrustworthy central database operator. In recent years, Goolge [Erlingsson et al., 2014], Apple [Thakurta et al., 2017] and Microsoft [Ding et al., 2017] have all developed and deployed local differential privacy procedures in data collection.

There are two type of possible unscrupulous access to the raw data $\boldsymbol{X}_1$ to be addressed. The first is that the data collector is untrustworthy. The second is that an unscrupulous party might break in to the server containing data collected by an honest data collector. In the differential privacy literature, the first type is handled by using local differential privacy procedures, while the second type is addressed via pan-private data analysis [Dwork et al., 2010]. Our TM^2 scheme protects against both type of unscrupulous accesses, but only allow for a one-shot collection for each individual's data.

While both the local differential privacy procedures and the TM^2 scheme can provide protection against unscrupulous accesses, the goals are somewhat different. The TM^2 scheme aims to collect a masked data set that preserves the first two statistical moments of the variables (note that $\boldsymbol{X}_1^T\boldsymbol{X}_1$ is knowable from the publicly available $\boldsymbol{AX}_1$). This allows exact statistical inferences on quantities depending on these statistical moments. The local differential privacy methods, on the other hand, aims to provide a stronger privacy protection under composition of multiple data collections/accesses.

The idea of the TM^2 scheme is similar to secure multi-party computation (SMC) procedures, in that this scheme tries to distribute information among parties so that each party does not get access to individual level data other than its own. There are also important differences between TM^2 and SMC. They differ in their designed purposes even though both want parties to cooperate in a joint task while keeping privacy. SMC is designed to conduct joint statistical analysis without the parties revealing their data to each other. TM^2 wants to collect the masked data set, which enables statistical analysis, without parties revealing the actual data to the data collector. Operationally, SMC requires distributed storage of data as well as distributed computation. Specifically, if we require that the private data of parties never leave their devices, then SMC needs the parties to stand by ready for any statistical analysis that may occur much later in the future. In contrast, the TM^2 method is only distributive in the data collection stage. The private data leaves the parties' devices in a masked form, and later is centrally stored in masked form $\boldsymbol{A}\boldsymbol{X}_1$. Since all future statistical analysis is conducted on the publicly released $\boldsymbol{A}\boldsymbol{X}_1$, there is no need for the parties to be available for future analysis.

In this paper, we presented a privacy analysis clearly separating the risks coming from support restriction and the risks of probabilistic attacks beyond the support restriction. With the analysis, we show that the TM^2 scheme is safe to collect a synthetic data set $\boldsymbol{A}\boldsymbol{X}_1$ which is a random orthogonal transformation of the raw data set $\boldsymbol{X}_1$. All information during the data collection procedure is masked, and no one during the procedure can access the raw data set. This removes the issue of trusting a data record keeper and provides a new tool for researchers to collect data allowing exact statistical inference for linear models while provide a privacy protection: no hacking attack against a party in the data collection procedure can access real individual level data since all parties do not have enough information to infer the private individual data.

Acknowledgment

Research reported in this publication was supported by the National Institute Of General Medical Sciences of the National Institutes of Health under Award Number R01GM118737, and by National Science Foundation under Award Number CNS-1563697.

References

T. W. Anderson, I. Olkin, and L. G. Underhill. Generation of random orthogonal matrices. *SIAM Journal on Scientific and Statistical Computing*, 8(4):625–629, 1987. doi: 10.1137/0908055. URL https://doi.org/10.1137/0908055.

G. Blair, K. Imai, and Y. Zhou. Design and analysis of the randomized response technique. *Journal of the American Statistical Association*, 110(511):1304–1319, 2015.

R. Brand. Microdata protection through noise addition. In *Inference control in statistical databases*, pages 97–116. Springer, 2002.

J. Burridge. Information preserving statistical obfuscation. *Statistics and Computing*, 13(4):321–327, 2003.

G. Cormode, S. Jha, T. Kulkarni, N. Li, D. Srivastava, and T. Wang. Privacy at scale: Local differential privacy in practice. In *Proceedings of the 2018 International Conference on Management of Data*, pages 1655–1658. ACM, 2018.

B. Ding, J. Kulkarni, and S. Yekhanin. Collecting telemetry data privately. In *Advances in Neural Information Processing Systems*, pages 3571–3580, 2017.

J. Drechsler and J. P. Reiter. Sampling with synthesis: A new approach for releasing public use census microdata. *Journal of the American Statistical Association*, 105(492): 1347–1357, 2010.

J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Minimax optimal procedures for locally private estimation. *Journal of the American Statistical Association*, (accepted), 2017.

C. Dwork. Differential privacy in M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, eds., ICALP (2). *Lecture Notes in Computer Science*, 4052:1–12, 2006.

C. Dwork and M. Naor. On the difficulties of disclosure prevention in statistical databases or the case for differential privacy. *Journal of Privacy and Confidentiality*, 2(1):8, 2008.

C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Third Conference on Theory of Cryptography*, TCC'06, pages 265–284, Berlin, Heidelberg, 2006. Springer-Verlag. ISBN 3-540-32731-2, 978-3-540-32731-8. doi: 10.1007/11681878_14. URL http://dx.doi.org/10.1007/11681878_14.

C. Dwork, M. Naor, T. Pitassi, G. N. Rothblum, and S. Yekhanin. Pan-private streaming algorithms. In *ICS*, pages 66–80, 2010.

Ú. Erlingsson, V. Pihur, and A. Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, pages 1054–1067. ACM, 2014.

A. Evfimievski, J. Gehrke, and R. Srikant. Limiting privacy breaches in privacy preserving data mining. In *Proceedings of the Twenty-second ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, PODS '03, pages 211–222, New York, NY, USA, 2003. ACM. ISBN 1-58113-670-6. doi: 10.1145/773153.773174. URL http://doi.acm.org/10.1145/773153.773174.

B. C. M. Fung, K. Wang, R. Chen, and P. S. Yu. Privacy-preserving data publishing: A survey of recent developments. *ACM Comput. Surv.*, 42(4):14:1–14:53, June 2010. ISSN 0360-0300. doi: 10.1145/1749603.1749605. URL http://doi.acm.org/10.1145/1749603.1749605.

R. M. Heiberger. Algorithm as 127: Generation of random orthogonal matrices. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 27(2):199–206, 1978. ISSN 00359254, 14679876. URL http://www.jstor.org/stable/2346957.

Huffington Post. *Citigroup: $2.7 Million Stolen From Customers As Result Of Hacking. http://www.huffingtonpost.com/2011/06/27/citigroup-hack_n_885045.html*. 2011.

S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith. What can we learn privately? *SIAM J. Comput.*, 40(3):793–826, June 2011. ISSN 0097-5397. doi: 10.1137/090756090. URL http://dx.doi.org/10.1137/090756090.

M. Ledoux, B. Rider, et al. Small deviations for beta ensembles. *Electronic Journal of Probability*, 15:1319–1343, 2010.

N. Li, T. Li, and S. Venkatasubramanian. t-closeness: Privacy beyond k-anonymity and l-diversity. In *Data Engineering, 2007. ICDE 2007. IEEE 23rd International Conference on*, pages 106–115. IEEE, 2007.

K. Liu, H. Kargupta, and J. Ryan. Random projection-based multiplicative data perturbation for privacy preserving distributed data mining. *IEEE Transactions on knowledge and Data Engineering*, 18(1):92–106, 2006.

A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkitasubramaniam. L-diversity: Privacy beyond k-anonymity. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 1 (1):3, 2007.

Reuters. *Target To Pay $10 Million To Settle Lawsuit From Massive Data Breach. http://www.huffingtonpost.com/2015/03/18/target-hack-settlement_n_6899290.html.* 2015.

Reuters. *Equifax Says Hack Potentially Exposed Details Of 143 Million Consumers. https://www.huffingtonpost.com/entry/equifax-says-hack-potentially-exposed-details-of-143-million-consumers_us_59b1bc2de4b0354e4410b33e.* 2017.

D. B. Rubin. Satisfying confidentiality constraints through the use of synthetic multiply imputed microdata. *Journal of Official Statistics*, 9:461–468, 1993.

L. Sweeney. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(05):557–570, 2002.

A. G. Thakurta, A. H. Vyrros, U. S. Vaishampayan, G. Kapoor, J. Freudiger, V. R. Sridhar, and D. Davidson. Learning new words. Number US9594741B1. US Patent 9594741B1, Mar 2017.

D. Ting, S. E. Fienberg, and M. Trottini. Random orthogonal matrix masking methodology for microdata release. *International Journal of Information and Computer Securityroke*, 2 (1):86–105, 2008.

S. L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.

S. S. Wu, S. Chen, A. Bhattacharjee, and Y. He. Collusion resistant multi-matrix masking for privacy-preserving data collection. In *IEEE 3rd International Conference on Big Data Security on Cloud (BigDataSecurity)*, pages 1–7. IEEE, 2017a.

S. S. Wu, S. Chen, D. L. Burr, and L. Zhang. A new data collection technique for preserving privacy. *Journal of Privacy and Confidentiality*, 7(3):99–129, 2017b.

L. Zhang. *On security properties of Random Matrix Masking.* PhD thesis, University of Florida, January 2014. URL http://search.proquest.com/docview/1876889174/.

Appendix A. A Counter Example

We illustrate that $p \geq n$ and the full rank condition on $\boldsymbol{X}$ are needed for the privacy preservation in the TM^2 scheme through a simple counter example here. We consider a 3×2 matrix $\boldsymbol{X}$, where the first column $\boldsymbol{X}_1$ contains binary sensitive information and the second column $\boldsymbol{X}_2$ contains continuous random noise. Suppose that only one of the three individuals answered "1" on the sensitive question, so that the data matrix is

$$\boldsymbol{X} = \begin{pmatrix} x_{11} & x_{12} \\ x_{21} & x_{22} \\ x_{31} & x_{32} \end{pmatrix} = \begin{pmatrix} 1 & a_1 \\ 0 & a_2 \\ 0 & a_3 \end{pmatrix}. \tag{A.1}$$

We decompose $\boldsymbol{X} = \begin{pmatrix} \boldsymbol{X}_a \\ \boldsymbol{X}_b \end{pmatrix}$ with the first two rows as $\boldsymbol{X}_a$ and the last row as $\boldsymbol{X}_b$. Without loss of generality, we assume that $a_2 \neq 0$ so that $\begin{pmatrix} 1 & a_1 \\ 0 & a_2 \end{pmatrix}$ is non-singular, and we assume that a_3/a_2 is not an integer. Then the first column $\boldsymbol{X}_1 = (1, 0, 0)^T$ can be uniquely determined from the masked data $\boldsymbol{M}_R$.

To see this, we decompose $\boldsymbol{M}_R = \begin{pmatrix} \boldsymbol{M}_a \\ \boldsymbol{M}_b \end{pmatrix}$ similarly as in the decomposition of $\boldsymbol{X} = \begin{pmatrix} \boldsymbol{X}_a \\ \boldsymbol{X}_b \end{pmatrix}$. Then $\boldsymbol{M}_b \boldsymbol{M}_a^{-1} = (\boldsymbol{X}_b \boldsymbol{B})(\boldsymbol{X}_a \boldsymbol{B})^{-1} = \boldsymbol{X}_b \boldsymbol{B}\boldsymbol{B}^{-1} \boldsymbol{X}_a^{-1} = \boldsymbol{X}_b \boldsymbol{X}_a^{-1}$ is known to anyone with access to $\boldsymbol{M}_R$. Using (A.1), this means $\boldsymbol{X}_b \boldsymbol{X}_a^{-1} = (0, a_3/a_2)$ is determined from the masked data $\boldsymbol{M}_R$. Then the first element in $(\boldsymbol{X}_b\boldsymbol{X}_a^{-1})\boldsymbol{X}_a = \boldsymbol{X}_b$ indicates that $(0, a_3/a_2)(x_{11}, x_{21})^T = x_{31}$. That is, $(a_3/a_2)x_{21} = x_{31}$.

Since x_{21} and x_{31} are binary entries in $\boldsymbol{X}_1$ and a_3/a_2 is not an integer, the attacker can infer from $(a_3/a_2)x_{21} = x_{31}$ that $x_{21} = x_{31} = 0$. Then we must have $x_{11} = 1$ due to $\boldsymbol{M}_R$ (and thus $\boldsymbol{X}$) being full-rank. That is, we now know every entry in $\boldsymbol{X}_1 = (1, 0, 0)^T$ from the masked data $\boldsymbol{M}_R$.

Notice that according to Lemma 3.3, a strong collection obfuscating procedure would not have allowed this identification of individual data from the random permutation. There is indeed additional privacy loss without assuming $p < n$. In general, when $p < n$ and $\boldsymbol{M}_R$ is full rank, $(\boldsymbol{M}_b\boldsymbol{M}_a^{-1})\boldsymbol{X}_a = \boldsymbol{X}_b$ along with knowledge of the data type may leak sensitive information about original data $\boldsymbol{X}$.

Appendix B. Proof of Theorem 3.7.

Proof. Same arguments in proof of Theorem 3.6 shows that (3.3) holds. Therefore we only need to show that there exist $\tilde{\boldsymbol{X}}_2$ and $\tilde{\boldsymbol{B}}$ satisfying (3.4): $(\boldsymbol{P}\boldsymbol{X}_1, \tilde{\boldsymbol{U}})\tilde{\boldsymbol{B}} = \boldsymbol{X}\boldsymbol{B} = \boldsymbol{M}_R$.

Using condition (3.5), we have

$$\begin{aligned}\lambda_{min}(\boldsymbol{X}_1\boldsymbol{X}_1^T + \boldsymbol{X}_2\boldsymbol{X}_2^T - \boldsymbol{P}\boldsymbol{X}_1\boldsymbol{X}_1^T\boldsymbol{P}^T) &\geq \lambda_{min}(\boldsymbol{X}_1\boldsymbol{X}_1^T) + \lambda_{min}(\boldsymbol{X}_2\boldsymbol{X}_2^T) - \lambda_{max}(\boldsymbol{P}\boldsymbol{X}_1\boldsymbol{X}_1^T\boldsymbol{P}^T)\\ &\geq \lambda_{min}(\boldsymbol{X}_2\boldsymbol{X}_2^T) - \lambda_{max}(\boldsymbol{X}_1\boldsymbol{X}_1^T)\\ &> 0.\end{aligned}$$

Hence $\boldsymbol{X}_1\boldsymbol{X}_1^T + \boldsymbol{X}_2\boldsymbol{X}_2^T - \boldsymbol{P}\boldsymbol{X}_1\boldsymbol{X}_1^T\boldsymbol{P}^T$ is a positive definite matrix. Therefore, there exists a matrix $\tilde{\boldsymbol{U}}$ such that

$$\tilde{\boldsymbol{U}}\tilde{\boldsymbol{U}}^T = \boldsymbol{X}_1\boldsymbol{X}_1^T + \boldsymbol{X}_2\boldsymbol{X}_2^T - \boldsymbol{P}\boldsymbol{X}_1\boldsymbol{X}_1^T\boldsymbol{P}^T. \tag{B.1}$$

This is equivalent to

$$(\boldsymbol{P}\boldsymbol{X}_1, \tilde{\boldsymbol{U}})(\boldsymbol{P}\boldsymbol{X}_1, \tilde{\boldsymbol{U}})^T = \tilde{\boldsymbol{U}}\tilde{\boldsymbol{U}}^T + \boldsymbol{P}\boldsymbol{X}_1\boldsymbol{X}_1^T\boldsymbol{P}^T = \boldsymbol{X}_1\boldsymbol{X}_1^T + \boldsymbol{X}_2\boldsymbol{X}_2^T = \boldsymbol{X}\boldsymbol{X}^T = \boldsymbol{M}_R\boldsymbol{M}_R^T. \tag{B.2}$$

Now we apply a singular value decomposition on $\boldsymbol{M}_R = \boldsymbol{S}\boldsymbol{D}\boldsymbol{V}$ where $\boldsymbol{S} \in \mathcal{O}_n$, $\boldsymbol{V} \in \mathcal{O}_p$ and $\boldsymbol{D}$ is a diagonal matrix with nonincreasing nonnegative diagonal elements. Then, due to (B.2), the singular decomposition of $(\boldsymbol{P}\boldsymbol{X}_1, \tilde{\boldsymbol{U}})$ is $\boldsymbol{S}\boldsymbol{D}\tilde{\boldsymbol{V}}$ for a $\tilde{\boldsymbol{V}} \in \mathcal{O}_p$. Therefore $\tilde{\boldsymbol{B}} = \tilde{\boldsymbol{V}}^T\boldsymbol{V}$ is the orthogonal matrix satisfies (3.4).

□

Appendix C. Bound on condition (3.7).

To achieve the condition (3.7) we can specify the noise distribution to have large noise values. Let x_{max} denote the largest possible absolute value of entries in $\boldsymbol{X}_1$. Then

$$\lambda_{max}(\boldsymbol{X}_1\boldsymbol{X}_1^T) = \|\boldsymbol{X}_1\|_2^2 \leq \|\boldsymbol{X}_1\|_F^2 = \sum_{i=1}^{n}\sum_{j=1}^{p_1} x_{1,ij}^2 \leq np_1x_{max}^2 = C_n,$$

where $\|\cdot\|_2$ and $\|\cdot\|_F$ are the operator norm and Frobenius norm respectively. Note that $C_n = np_1x_{max}^2$ is a known constant to the designer of the TM^2 scheme. Condition (3.5) holds when $\lambda_{min}(\boldsymbol{X}_2\boldsymbol{X}_2^T)$ exceeds this constant.

We require each data provider to generate a p_2-dimensional random noise vector with i.i.d. Gaussian elements of mean zero and variance σ^2. Assume that $\gamma = p_2/n > 1$, when $n \to \infty$, Corollary 13 in Ledoux et al. [2010] provides a probability bound on the $\lambda_{min}(\boldsymbol{X}_2\boldsymbol{X}_2^T)$ for any $\delta > 0$:

$$Pr[\lambda_{min}(\boldsymbol{X}_2\boldsymbol{X}_2^T) \leq (\sqrt{\gamma}-1)^2 n\sigma^2(1-\delta)] \leq C_0 e^{-n\delta^{3/2}/C_0},$$

for some constant C_0. The bound on the right side decrease exponentially in n so that, for large n, it can be made smaller than ϵ for a $\delta < 1$. Choosing $\sigma^2 > C_n/[(\sqrt{\gamma}-1)^2 n(1-\delta)]$ will ensure that (3.7) holds.

Appendix D. Proof of Theorem 3.8.

Proof. We study the posterior density

$$\pi_{\boldsymbol{X}_1|(\boldsymbol{M}_L,\boldsymbol{M}_R)}(\boldsymbol{x}_1|\boldsymbol{m}_L,\boldsymbol{m}_R) = \frac{\pi_{(\boldsymbol{X}_1,\boldsymbol{M}_L,\boldsymbol{M}_R)}(\boldsymbol{x}_1,\boldsymbol{m}_L,\boldsymbol{m}_R)}{\int\limits_{\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{m}_L,\boldsymbol{m}_R)} \pi_{(\boldsymbol{X}_1,\boldsymbol{M}_L,\boldsymbol{M}_R)}(\boldsymbol{x}_1^*,\boldsymbol{m}_L,\boldsymbol{m}_R)d\boldsymbol{x}_1^*}, \tag{D.1}$$

and compare it with the prior density $\pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1)$ restricted on the support $\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{M}_L,\boldsymbol{M}_R)$.

Recall that the probability densities for $\boldsymbol{X}_1$, $\boldsymbol{X}_2$, $\boldsymbol{A}$ and $\boldsymbol{B}$ at values $\boldsymbol{X}_1 = \boldsymbol{x}_1$, $\boldsymbol{X}_2 = \boldsymbol{x}_2$, $\boldsymbol{A} = \boldsymbol{a}$ and $\boldsymbol{B} = \boldsymbol{b}$ are denoted respectively as $\pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1)$, $\pi_{\boldsymbol{X}_2}(\boldsymbol{x}_2)$, $\pi_{\boldsymbol{A}}(\boldsymbol{a})$ and $\pi_{\boldsymbol{B}}(\boldsymbol{b})$. Due to the independence of the generation mechanism of these quantities, their joint density is

$$\pi_{(\boldsymbol{X}_1,\boldsymbol{X}_2,\boldsymbol{A},\boldsymbol{B})}(\boldsymbol{x}_1,\boldsymbol{x}_2,\boldsymbol{a},\boldsymbol{b}) = \pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1)\pi_{\boldsymbol{X}_2}(\boldsymbol{x}_2)\pi_{\boldsymbol{A}}(\boldsymbol{a})\pi_{\boldsymbol{B}}(\boldsymbol{b}), \tag{D.2}$$

for $(\boldsymbol{X}_1,\boldsymbol{X}_2,\boldsymbol{A},\boldsymbol{B}) \in \mathcal{S}_{\boldsymbol{X}_1} \times \mathcal{S}_{\boldsymbol{X}_2} \times \mathcal{S}_{\boldsymbol{A}} \times \mathcal{S}_{\boldsymbol{B}}$.

Since the elements of $\boldsymbol{X}_2$ are i.i.d. from the Gaussian distribution $N(0,\sigma^2)$,

$$\pi_{\boldsymbol{X}_2}(\boldsymbol{x}_2) = \frac{1}{(\sqrt{2\pi}\sigma)^{np_2}} e^{-\frac{\sum_{1\leq i\leq n,1\leq j\leq p_2} x_{2,ij}^2}{2\sigma^2}} = f(\|\boldsymbol{x}_2\|_F^2),$$

where $f(x) = \frac{1}{(\sqrt{2\pi}\sigma)^{np_2}} e^{-\frac{x}{2\sigma^2}}$ and $\|\boldsymbol{x}_2\|_F^2 = \sum_{1\leq i\leq n,1\leq j\leq p_2} x_{2,ij}^2$ with $\|\cdot\|_F$ denote the Frobenius norm. Thus the joint density becomes

$$\pi_{(\boldsymbol{X}_1,\boldsymbol{X}_2,\boldsymbol{A},\boldsymbol{B})}(\boldsymbol{x}_1,\boldsymbol{x}_2,\boldsymbol{a},\boldsymbol{b}) = \pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1)f(\|\boldsymbol{x}_2\|_F^2)\pi_{\boldsymbol{A}}(\boldsymbol{a})\pi_{\boldsymbol{B}}(\boldsymbol{b}). \tag{D.3}$$

We now plug (D.3) into (D.1) for calculation.

First, we calculate the numerator $\pi_{(\boldsymbol{X}_1,\boldsymbol{M}_L,\boldsymbol{M}_R)}(\boldsymbol{x}_1,\boldsymbol{m}_L,\boldsymbol{m}_R)$ in (D.1). We denote the restricted sample spaces of random variables $\boldsymbol{A}$ and $\boldsymbol{B}$ respectively given knowledge of some other quantities as:

$$\begin{aligned} \mathcal{S}_A(\boldsymbol{x}_1,\boldsymbol{m}_L) &= \{\boldsymbol{a} : \ \boldsymbol{a}\boldsymbol{x}_1 = \boldsymbol{m}_L\}, \\ \mathcal{S}_B(\boldsymbol{x}_1,\boldsymbol{m}_R) &= \{\boldsymbol{b} : \ \exists \boldsymbol{x}_2 \text{ such that } (\boldsymbol{x}_1,\boldsymbol{x}_2)\boldsymbol{b} = \boldsymbol{m}_R\}. \end{aligned} \tag{D.4}$$

Notice that given $\boldsymbol{A} = \boldsymbol{a}$ and $\boldsymbol{M}_L = \boldsymbol{m}_L$, then $\boldsymbol{X}_1 = \boldsymbol{a}^T\boldsymbol{m}_L$. Also, given $\boldsymbol{B} = \boldsymbol{b}$ and $\boldsymbol{M}_R = \boldsymbol{m}_R$, then $(\boldsymbol{X}_1,\boldsymbol{X}_2) = \boldsymbol{m}_R\boldsymbol{b}^T$ so that

$$\|\boldsymbol{X}_2\|_F^2 = trace(\boldsymbol{X}_2\boldsymbol{X}_2^T) = trace[\boldsymbol{m}_R\boldsymbol{b}^T\boldsymbol{b}\boldsymbol{m}_R^T - \boldsymbol{X}_1\boldsymbol{X}_1^T] = trace(\boldsymbol{m}_R\boldsymbol{m}_R^T) - trace(\boldsymbol{X}_1\boldsymbol{X}_1^T).$$

Hence given $\boldsymbol{A} = \boldsymbol{a}$, $\boldsymbol{M}_L = \boldsymbol{m}_L$ and $\boldsymbol{B} = \boldsymbol{b}$, we have

$$\|\boldsymbol{X}_2\|_F^2 = trace(\boldsymbol{m}_R\boldsymbol{m}_R^T) - trace(\boldsymbol{a}^T\boldsymbol{m}_L\boldsymbol{m}_L^T\boldsymbol{a}) = trace(\boldsymbol{m}_R\boldsymbol{m}_R^T) - trace(\boldsymbol{m}_L\boldsymbol{m}_L^T).$$

Then using this and equation (D.3), we have

$$\begin{aligned}
& \pi_{(\boldsymbol{X}_1,\boldsymbol{M}_L,\boldsymbol{M}_R)}(\boldsymbol{x}_1,\boldsymbol{m}_L,\boldsymbol{m}_R) \\
= & \int\limits_{\mathcal{S}_A(\boldsymbol{x}_1,\boldsymbol{m}_L)} \{ \int\limits_{\mathcal{S}_B(\boldsymbol{x}_1,\boldsymbol{m}_R)} \pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1) f[trace(\boldsymbol{m}_R\boldsymbol{m}_R^T) - trace(\boldsymbol{m}_L\boldsymbol{m}_L^T)] \pi_{\boldsymbol{A}}(\boldsymbol{a})\pi_{\boldsymbol{B}}(\boldsymbol{b}) d\boldsymbol{b}\} d\boldsymbol{a} \\
= & \pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1) f[trace(\boldsymbol{m}_R\boldsymbol{m}_R^T) - trace(\boldsymbol{m}_L\boldsymbol{m}_L^T)] \int\limits_{\mathcal{S}_A(\boldsymbol{x}_1,\boldsymbol{m}_L)} \{ \int\limits_{\mathcal{S}_B(\boldsymbol{x}_1,\boldsymbol{m}_R)} \pi_{\boldsymbol{A}}(\boldsymbol{a})\pi_{\boldsymbol{B}}(\boldsymbol{b}) d\boldsymbol{b}\} d\boldsymbol{a} \\
= & \pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1) f[trace(\boldsymbol{m}_R\boldsymbol{m}_R^T) - trace(\boldsymbol{m}_L\boldsymbol{m}_L^T)] \int\limits_{\mathcal{S}_A(\boldsymbol{x}_1,\boldsymbol{m}_L)} [\int\limits_{\mathcal{S}_B(\boldsymbol{x}_1,\boldsymbol{m}_R)} \pi_{\boldsymbol{B}}(\boldsymbol{b}) d\boldsymbol{b}] \pi_{\boldsymbol{A}}(\boldsymbol{a}) d\boldsymbol{a} \\
= & \pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1) f[trace(\boldsymbol{m}_R\boldsymbol{m}_R^T) - trace(\boldsymbol{m}_L\boldsymbol{m}_L^T)] [\int\limits_{\mathcal{S}_A(\boldsymbol{x}_1,\boldsymbol{m}_L)} \pi_{\boldsymbol{A}}(\boldsymbol{a}) d\boldsymbol{a}] [\int\limits_{\mathcal{S}_B(\boldsymbol{x}_1,\boldsymbol{m}_R)} \pi_{\boldsymbol{B}}(\boldsymbol{b}) d\boldsymbol{b}].
\end{aligned} \tag{D.5}$$

Now for any pair of $\boldsymbol{x}_1$ and $\boldsymbol{x}_1^*$ that both belongs to $\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{m}_L,\boldsymbol{m}_R)$, there exist $(\boldsymbol{a},\boldsymbol{a}^*)$ such that $\boldsymbol{a}\boldsymbol{x}_1 = \boldsymbol{m}_L = \boldsymbol{a}^*\boldsymbol{x}_1^*$. Denote $\boldsymbol{A}_0 = (\boldsymbol{a})^{-1}\boldsymbol{a}^*$. Then $\boldsymbol{A}_0\boldsymbol{x}_1^* = (\boldsymbol{a})^{-1}\boldsymbol{m}_L = \boldsymbol{x}_1$ and $\boldsymbol{A}_0^{-1}\boldsymbol{x}_1 = \boldsymbol{x}_1^*$. Hence for any $\tilde{\boldsymbol{a}} \in \mathcal{S}_A(\boldsymbol{x}_1,\boldsymbol{m}_L)$ we have

$$\tilde{\boldsymbol{a}}\boldsymbol{A}_0\boldsymbol{x}_1^* = \tilde{\boldsymbol{a}}\boldsymbol{x}_1 = \boldsymbol{m}_L,$$

i.e., $\tilde{\boldsymbol{a}}\boldsymbol{A}_0 \in \mathcal{S}_A(\boldsymbol{x}_1^*,\boldsymbol{m}_L)$. On the other hand, for any $\bar{\boldsymbol{a}} \in \mathcal{S}_A(\boldsymbol{x}_1^*,\boldsymbol{m}_L)$, $\bar{\boldsymbol{a}}\boldsymbol{A}_0^{-1}\boldsymbol{x}_1 = \bar{\boldsymbol{a}}\boldsymbol{x}_1^* = \boldsymbol{m}_L$, i.e., $\bar{\boldsymbol{a}}\boldsymbol{A}_0^{-1} \in \mathcal{S}_A(\boldsymbol{x}_1,\boldsymbol{m}_L)$. Taken together, we have a one-to-one mapping between the two sets $\mathcal{S}_A(\boldsymbol{x}_1,\boldsymbol{m}_L)$ and $\mathcal{S}_A(\boldsymbol{x}_1^*,\boldsymbol{m}_L)$. Particularly,

$$\mathcal{S}_A(\boldsymbol{x}_1^*,\boldsymbol{m}_L) = \mathcal{S}_A(\boldsymbol{x}_1,\boldsymbol{m}_L)\boldsymbol{A}_0. \tag{D.6}$$

Hence for uniform density $\pi_{\boldsymbol{A}} = \pi_0$, (D.6) and Lemma 3.4 implies that

$$\int\limits_{\mathcal{S}_A(\boldsymbol{x}_1,\boldsymbol{m}_L)} \pi_{\boldsymbol{A}}(\boldsymbol{a}) d\boldsymbol{a} = \int\limits_{\mathcal{S}_A(\boldsymbol{x}_1^*,\boldsymbol{m}_L)} \pi_{\boldsymbol{A}}(\boldsymbol{a}) d\boldsymbol{a}. \tag{D.7}$$

Plug (D.5) and (D.7) into (D.1) and cancel the common factors, we get

$$\begin{aligned}
& \pi_{\boldsymbol{X}_1|(\boldsymbol{M}_L,\boldsymbol{M}_R)}(\boldsymbol{x}_1|\boldsymbol{m}_L,\boldsymbol{m}_R) \\
= & \frac{\pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1)[\int\limits_{\mathcal{S}_B(\boldsymbol{x}_1,\boldsymbol{m}_R)} \pi_{\boldsymbol{B}}(\boldsymbol{b})d\boldsymbol{b}]}{\int\limits_{\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{m}_L,\boldsymbol{m}_R)} \pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1^*)[\int\limits_{\mathcal{S}_B(\boldsymbol{x}_1^*,\boldsymbol{m}_R)} \pi_{\boldsymbol{B}}(\boldsymbol{b})d\boldsymbol{b}]d\boldsymbol{x}_1^*},
\end{aligned} \tag{D.8}$$

Next, for any pair of $\boldsymbol{x}_1$ and $\boldsymbol{x}_1^*$ that both belongs to $\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{m}_L,\boldsymbol{m}_R)$, there exist $(\boldsymbol{x}_2,\boldsymbol{b})$ and $(\boldsymbol{x}_2^*,\boldsymbol{b}^*)$ such that $(\boldsymbol{x}_1,\boldsymbol{x}_2)\boldsymbol{b} = \boldsymbol{m}_R = (\boldsymbol{x}_1^*,\boldsymbol{x}_2^*)\boldsymbol{b}^*$. Let $\boldsymbol{B}_0 = \boldsymbol{b}(\boldsymbol{b}^*)^{-1}$. Then $(\boldsymbol{x}_1,\boldsymbol{x}_2)\boldsymbol{B}_0 = (\boldsymbol{x}_1^*,\boldsymbol{x}_2^*)$. Similar to (D.6), we have

$$\mathcal{S}_B(\boldsymbol{x}_1,\boldsymbol{m}_R) = \boldsymbol{B}_0\mathcal{S}_B(\boldsymbol{x}_1^*,\boldsymbol{m}_R). \tag{D.9}$$

For uniform density $\pi_{\boldsymbol{B}} = \pi_0$, using (D.9), Lemma 3.4 implies that

$$\int\limits_{\mathcal{S}_B(\boldsymbol{x}_1,\boldsymbol{m}_R)} \pi_{\boldsymbol{B}}(\boldsymbol{b}) d\boldsymbol{b} = \int\limits_{\mathcal{S}_B(\boldsymbol{x}_1^*,\boldsymbol{m}_R)} \pi_{\boldsymbol{B}}(\boldsymbol{b}) d\boldsymbol{b} \tag{D.10}$$

Plug-in (D.10) into equation (D.8), we have

$$\pi_{\boldsymbol{X}_1|(\boldsymbol{M}_L,\boldsymbol{M}_R)}(\boldsymbol{x}_1|\boldsymbol{m}_L,\boldsymbol{m}_R) = \frac{\pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1)}{\int\limits_{\mathcal{S}_{\boldsymbol{X}_1}(\boldsymbol{m}_L,\boldsymbol{m}_R)} \pi_{\boldsymbol{X}_1}(\boldsymbol{x}_1^*)d\boldsymbol{x}_1^*}.$$

□