
An Optimal and Scalable Matrix Mechanism for Noisy Marginals under Convex Loss Functions

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Noisy marginals are a common form of confidentiality-protecting data release and
2 are useful for many downstream tasks such as contingency table analysis, construc-
3 tion of Bayesian networks, and even synthetic data generation. Privacy mechanisms
4 that provide unbiased noisy answers to linear queries (such as marginals) are known
5 as matrix mechanisms.
6 We propose ResidualPlanner, a matrix mechanism for marginals with Gaussian
7 noise that is both optimal and scalable. ResidualPlanner can optimize for many
8 loss functions that can be written as a convex function of marginal variances (prior
9 work was restricted to just one predefined objective function). ResidualPlanner can
10 optimize the accuracy of marginals in large scale settings in seconds, even when the
11 previous state of the art (HDMM) runs out of memory. It even runs on datasets with
12 100 attributes in a couple of minutes. Furthermore ResidualPlanner can efficiently
13 compute variance/covariance values for each marginal (prior methods quickly run
14 out of memory, even for relatively small datasets).

15 1 Introduction

16 Marginals are tables of counts on a set of attributes (e.g., how many people there are for each
17 combination of race and gender). They are one of the most common formats for the dissemination of
18 statistical data [8, 2], studying correlations between attributes, and are sufficient statistics for loglinear
19 models, including Bayesian networks and Markov random fields. For this reason, a lot of work in
20 the differential privacy literature has considered how to produce a set of noisy marginals that is both
21 privacy-preserving and accurate.

22 One line of work, called the *matrix mechanism* [32, 52, 30, 53, 37, 51, 46, 18, 42] designs algorithms
23 for answering linear queries (such as marginals) so that the privacy-preserving noisy answers are
24 accurate, unbiased, and have a simple distribution (e.g., multivariate normal). These crucial properties
25 allow statisticians to work with the data, model error due to data collection (sampling error) and
26 error due to privacy protections. It enables valid confidence intervals and hypothesis tests and
27 other methods for quantifying the uncertainty of a statistical analysis (e.g., [20, 29, 50, 25, 26]).
28 Incidentally, sets of noisy marginals are also used to generate differentially private synthetic data
29 (e.g., [54, 4, 41, 10]).

30 For the case of marginals, significant effort has been spent in designing optimal or nearly optimal
31 matrix mechanisms for just a single objective function (total variance of all the desired marginals)
32 [32, 49, 13, 52, 53, 31, 51] and each new objective function requires significant additional effort
33 [6, 18, 42, 46]. However, existing optimal solutions do not scale and additional effort is needed to
34 design scalable, but suboptimal, matrix mechanisms for marginals [37, 38]. Furthermore, computing
35 the individual variances of the desired noisy marginals is a slow process and more difficult is
36 computing the covariance between cells in the same marginal.

37 **Contributions.** Our paper addresses these problems with a novel matrix mechanisms called ResidualPlanner. It can optimize for a wide variety of convex objective functions and return solutions
38 that are guaranteed to be optimal under Gaussian noise. It is highly scalable – running in seconds
39 even when other scalable algorithms run out of memory. It also efficiently returns the variance and
40 covariances of each cell of the desired marginals. It leverages the following insights. Since a dataset
41 can be represented as a vector \mathbf{x} of counts, and since a marginal query on a set \mathbf{A} of attributes can be
42 represented as a matrix $\mathbf{Q}_{\mathbf{A}}$ (with $\mathbf{Q}_{\mathbf{A}}\mathbf{x}$ being the true answer to the marginal query), we find a new
43 linearly independent basis that can parsimoniously represent both a marginal $\mathbf{Q}_{\mathbf{A}}$ and the “difference”
44 between two marginals $\mathbf{Q}_{\mathbf{A}}$ and $\mathbf{Q}_{\mathbf{A}'}$ (subspace spanned by the rows of $\mathbf{Q}_{\mathbf{A}}$ that is orthogonal to the
45 rows of $\mathbf{Q}_{\mathbf{A}'}$). Using parsimonious linear bases, instead of overparametrized mechanisms, accounts
46 for the scalability. Optimality results from a deep analysis of the symmetry that marginals impose on
47 the optimal solution – the same linear basis is optimal for a wide variety of loss functions.

49 2 Preliminaries

50 A dataset $\mathcal{D} = \{r_1, \dots, r_n\}$ is a collection of records. Each record r_i contains n_a attributes
51 Att_1, \dots, Att_{n_a} and each attribute Att_j can take values $a_1^{(j)}, \dots, a_{|Att_j|}^{(j)}$. An attribute value $a_i^{(j)}$ for
52 attribute Att_j can be represented as a vector using one-hot encoding. Specifically, let $e_i^{(j)}$ be a row
53 vector of size $|Att_j|$ with a 1 in component i and 0 everywhere else. In this way $e_i^{(j)}$ represents the
54 attribute value $a_i^{(j)}$. A record r with attributes $Att_1 = a_{i_1}^{(1)}, Att_2 = a_{i_2}^{(2)}, \dots, Att_{n_a} = a_{i_{n_a}}^{(n_a)}$ can
55 thus be represented as the Kronecker product $e_{i_1}^{(1)} \otimes e_{i_2}^{(2)} \otimes \dots \otimes e_{i_{n_a}}^{(n_a)}$. This vector has a 1 in exactly
56 one position and 0s everywhere else. The position of the 1 is the *index* of record r . With this notation,
57 a dataset \mathcal{D} can be represented as a vector \mathbf{x} of integers. The value at index i is the number of times
58 the record associated with index i appears in \mathcal{D} . The number of components in this vector is denoted
59 as $d = \prod_{i=1}^{n_a} |Att_i|$. Given a subset \mathbf{A} of attributes, a *marginal query* on \mathbf{A} is a table of counts: for
60 each combination of values for the attributes in \mathbf{A} , it provides the number of records in \mathcal{D} having
61 those attribute value combinations. The marginal query can be represented as a Kronecker product
62 $\mathbf{Q}_{\mathbf{A}} = \mathbf{V}_1 \otimes \dots \otimes \mathbf{V}_{n_a}$ where \mathbf{V}_i is the row vector of all ones (i.e. $\mathbf{1}_{|Att_i|}^T$) if $Att_i \notin \mathbf{A}$ and \mathbf{V}_i is the
63 identity matrix $\mathcal{I}_{|Att_i|}$ if $Att_i \in \mathbf{A}$. The answer to the marginal query is obtained by evaluating the
64 matrix-vector product $\mathbf{Q}_{\mathbf{A}}\mathbf{x}$. For convenience, the notation introduced in this paper is summarized as
65 a table in the supplementary material.

66 **EXAMPLE 2.1.** As a running example, consider a dataset in which there are two attributes: Att_1
67 with values “yes” and “no”, and Att_2 with values “low”, “med”, “high”. The record (no, med)
68 is represented by the kron product $\begin{bmatrix} 0 & 1 \end{bmatrix} \otimes \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}$ and the marginal query on the set $\mathbf{A} = \{Att_1\}$ is
69 represented as $\mathbf{Q}_{\{Att_1\}} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \otimes \begin{bmatrix} 1 & 1 \end{bmatrix}$. Similarly, the marginal on attribute Att_2 is represented as
70 $\mathbf{Q}_{\{Att_2\}} = \begin{bmatrix} 1 & 1 \end{bmatrix} \otimes \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$. The query representing all one-way marginals is obtained by stacking
71 them: $\mathbf{Q}^{1\text{-way}} = \begin{bmatrix} \mathbf{Q}_{\{Att_1\}} \\ \mathbf{Q}_{\{Att_2\}} \end{bmatrix}$ and $\mathbf{Q}^{1\text{-way}}\mathbf{x}$ consists of the five query answers (number of records with
72 $Att_1 = \text{yes}$, number with $Att_1 = \text{no}$, number with $Att_2 = \text{low}$, etc.).

73 2.1 Differential Privacy

74 A mechanism \mathcal{M} is an algorithm whose input is a dataset and whose output provides privacy protec-
75 tions. Differential privacy is a family of privacy definitions that guide the behavior of mechanisms so
76 that they can inject enough noise to mask the effects of any individual. There are many versions of
77 differential privacy that support Gaussian noise, including approximate DP, zCDP, and Gaussian DP.

78 **DEFINITION 2.2 (Differential Privacy).** Let \mathcal{M} be a mechanism. For every pair of datasets $\mathcal{D}_1, \mathcal{D}_2$
79 that differ on the presence/absence of a single record and for all (measurable) sets $S \subseteq \text{range}(\mathcal{M})$,

- 80 • If $P(\mathcal{M}(\mathcal{D}_1) \in S) \leq e^\epsilon P(\mathcal{M}(\mathcal{D}_2) \in S) + \delta$ then \mathcal{M} satisfies (ϵ, δ) -approximate differential
81 privacy [17];
- 82 • If $\Phi^{-1}(P(\mathcal{M}(\mathcal{D}_1) \in S)) \leq \Phi^{-1}(P(\mathcal{M}(\mathcal{D}_2) \in S)) + \mu$, where Φ is the cdf of the standard
83 Gaussian distribution, then \mathcal{M} satisfies μ -Gaussian DP [15].
- 84 • If the Rényi divergence $D_\alpha(\mathcal{M}(\mathcal{D}_1) || \mathcal{M}(\mathcal{D}_2))$ between the output distributions of $\mathcal{M}(\mathcal{D}_1)$ and
85 $\mathcal{M}(\mathcal{D}_2)$ satisfies $D_\alpha(\mathcal{M}(\mathcal{D}_1) || \mathcal{M}(\mathcal{D}_2)) \leq \rho\alpha$ for all $\alpha > 1$, then \mathcal{M} satisfies ρ -zCDP [7].

86 Queries that are linear functions of the data vector \mathbf{x} can be answered privately using the *linear*
87 *Gaussian mechanism*, which adds correlated Gaussian noise to a linear function of \mathbf{x} , as follows.

88 DEFINITION 2.3 (Linear Gaussian Mechanism [46]). Given a $m \times d$ matrix \mathbf{B} and $m \times m$ covariance
89 matrix Σ , the (correlated) linear Gaussian mechanism \mathcal{M} is defined as $\mathcal{M}(\mathbf{x}) = \mathbf{B}\mathbf{x} + N(\mathbf{0}, \Sigma)$.
90 The privacy cost matrix of \mathcal{M} is defined as $\mathbf{B}^T \Sigma^{-1} \mathbf{B}$. The privacy cost of \mathcal{M} , denoted by
91 $pcost(\mathcal{M})$, is the largest diagonal of the privacy cost matrix and is used to compute the privacy
92 parameters: \mathcal{M} satisfies ρ -zCDP with $\rho = pcost(\mathcal{M})/2$ [46], satisfies (ϵ, δ) -approximate DP with
93 $\delta = \Phi(\sqrt{pcost(\mathcal{M})}/2 - \epsilon/\sqrt{pcost(\mathcal{M})}) - e^\epsilon \Phi(-\sqrt{pcost(\mathcal{M})}/2 - \epsilon/\sqrt{pcost(\mathcal{M})})$ (this is an
94 increasing function of $pcost(\mathcal{M})$ [5]), and satisfies μ -Gaussian DP with $\mu = \sqrt{pcost(\mathcal{M})}$ [15, 46].

95 The use of a non-identity covariance matrix will help simplify the description of the optimal choices
96 of \mathbf{B} and Σ . We note that an algorithm \mathcal{M}^* that releases the outputs of multiple linear Gaussian
97 mechanisms $\mathcal{M}_1, \dots, \mathcal{M}_k$ (with $\mathcal{M}_i(\mathbf{x}) = \mathbf{B}_i \mathbf{x} + N(\mathbf{0}, \Sigma_i)$) is again a linear Gaussian mechanism.
98 It is represented as $\mathcal{M}^*(\mathbf{x}) = \mathbf{B}^* \mathbf{x} + N(\mathbf{0}, \Sigma^*)$ with the matrix \mathbf{B}^* obtained by vertically stacking
99 the \mathbf{B}_i and with covariance Σ^* being a block-diagonal matrix where the blocks are the Σ_i . Its privacy
100 cost $pcost(\mathcal{M}^*) = pcost(\mathcal{M}_1, \dots, \mathcal{M}_k)$ is the largest diagonal entry of $\sum_{i=1}^k \mathbf{B}_i^T \Sigma_i^{-1} \mathbf{B}_i$.

101 2.2 Matrix Mechanism

102 The Matrix Mechanism [32, 52, 30, 53, 37, 38, 51, 46, 18, 42] is a framework for providing unbiased
103 privacy-preserving answers to a workload of linear queries, represented by a matrix \mathbf{W} (so that the
104 true non-private answer to the workload queries is $\mathbf{W}\mathbf{x}$). The matrix mechanism framework consists
105 of 3 steps: *select*, *measure*, and *reconstruct*. The purpose of the *select* phase is to determine *what* we
106 add noise to and *how much* noise to use. More formally, when a user’s preferred noise distribution is
107 Gaussian, the select phase chooses a Gaussian linear mechanism $\mathcal{M}(\mathbf{x}) \equiv \mathbf{B}\mathbf{x} + N(\mathbf{0}, \Sigma)$ whose
108 noisy output can be used to estimate the true query answer $\mathbf{W}\mathbf{x}$. Ideally, \mathcal{M} uses the least amount
109 of noise subject to privacy constraints (specified by a privacy definition and settings of its privacy
110 parameters). The *measure* phase runs the mechanism on the data to produce (noisy) privacy-preserving
111 outputs $\omega = \mathcal{M}(\mathbf{x})$. The *reconstruct* step uses ω to compute an unbiased estimate of $\mathbf{W}\mathbf{x}$. The
112 unbiased estimate is typically $\mathbf{W}(\mathbf{B}^T \Sigma^{-1} \mathbf{B})^\dagger \mathbf{B}^T \Sigma^{-1} \omega$, where \dagger represents the Moore-Penrose
113 pseudo-inverse. This is the best linear unbiased estimate of $\mathbf{W}\mathbf{x}$ that can be obtained from ω
114 [32]. This means that the goal of the select step is to optimize the choice of \mathbf{B} and Σ so that the
115 reconstructed answer is as accurate as possible, subject to privacy constraints. Ideally, a user would
116 specify their accuracy requirements using a loss function, but existing matrix mechanisms do not
117 allow this flexibility – they hard-code the loss function. The reason is each loss function requires
118 significant research and new optimization algorithms [53, 46, 18]. On top of this, existing optimal
119 matrix mechanism algorithms do not scale, while scalable matrix mechanisms are not guaranteed
120 to produce optimal solutions [37]. Additionally, the reconstruction phase should also compute the
121 variance of *each* workload answer. The variances are the diagonals of $\mathbf{W}(\mathbf{B}^T \Sigma^{-1} \mathbf{B})^\dagger \mathbf{W}^T$ and
122 making this computation scale is also challenging.

123 3 Additional Related Work

124 The marginal release mechanism by Barak et al. [6] predates the matrix mechanism [32, 52, 30, 53,
125 13, 43, 37, 51, 46, 18, 42, 38] and adds noise to the Fourier decomposition of marginals. We add
126 noise to a *different* basis, resulting in the scalability and optimality properties. The SVD bound [31]
127 is a lower bound on total matrix mechanism error when the loss function is the sum of variances. This
128 lower bound is tight for marginals and we use it as a sanity check for our results and implementation
129 (note ResidualPlanner provides optimal solutions even when the SVD bound is infeasible to compute).

130 Alternative approaches to the matrix mechanism can produce privacy preserving marginal query
131 answers that reduce variance by adding bias. This is often done by generating differentially private
132 synthetic data or other such data synopses from which marginals can be computed. State-of-the art
133 approaches iteratively ask queries and fit synthetic data to the resulting answers [22, 34, 4, 19, 39, 35,
134 44, 56]. For such mechanisms, it is difficult to estimate error of a query answer but recently AIM
135 [39] has made progress in upper bounding the error. PGM [41] provides a connection between the
136 matrix mechanism and this line of work, as it can postprocess noisy marginals into synthetic data. It

137 is a better alternative to sampling a synthetic dataset from models fit to carefully chosen marginals
 138 [54, 11, 55, 10]. Synthetic data for answering marginal queries can also be created from random
 139 projections [48], copulas [33, 3], and deep generative models [23, 1, 35].

140 With respect to the matrix mechanism, the reconstruction step is often one of the bottlenecks to
 141 scalability. While PGM [41] provides one solution, another proposal by McKenna et al. [40] is to
 142 further improve scalability by sacrificing some consistency (the answers to two different marginals
 143 may provide conflicting answers to submarginals they have in common). Work on differential privacy
 144 marginals has also seen extensions to hierarchical datasets, in which records form meaningful groups
 145 that need to be queried. That is, in addition to marginals on characteristics of people, marginals can
 146 be computed in different hierarchies such as geographic level (state, county, etc) and marginals on
 147 household composition (or other groupings of people) [2, 28, 36].

148 4 ResidualPlanner

149 ResidualPlanner is our proposed matrix mechanism for optimizing the accuracy of marginal queries
 150 with Gaussian noise. It is optimal and more scalable than existing approaches. It supports opti-
 151 mizing the accuracy of marginals under a wide variety of loss functions and provides exact vari-
 152 ances/covariances of the noisy marginals in closed-form. In this section, we first explain the loss
 153 functions it supports. We then describe the base mechanisms it uses to answer marginal queries. We
 154 next show how to reconstruct the marginal queries from the outputs of the base mechanisms and how
 155 to compute their variances in closed form. We then explain how to optimize these base mechanisms
 156 for different loss functions. The reason this selection step is presented last is because it depends on
 157 the closed form variance calculations. Then we analyze computational complexity.

158 4.1 Loss Functions Supported by ResidualPlanner

159 The loss functions we consider are based on the following principle: different marginals can have
 160 different relative importance but within a marginal, its cells are equally important. That is, a loss
 161 function can express that the two-way marginal on the attribute set {Race, Marital Status} is more
 162 important (i.e., requires more accuracy) than the 1-way marginal on {EducationLevel}, but all cells
 163 within the {Race, MaritalStatus} marginal are equally important. This is a commonly accepted
 164 principle for answering differentially private marginal queries (e.g., [32, 52, 30, 53, 37, 51, 46, 18,
 165 42, 39, 4, 34]) and is certainly true for the 2020 Census redistricting data [2].

166 Let $Wkload = \{\mathbf{A}_1, \dots, \mathbf{A}_k\}$ be a workload of marginals, where each \mathbf{A}_i is a subset of attributes
 167 and represents a marginal. E.g., $Wkload = \{\{\text{Race, MaritalStatus}\}, \{\text{EducationLevel}\}\}$ consists of
 168 2 marginals, a two-way marginal on Race/MaritalStatus, and a one-way marginal on Education. Let
 169 \mathcal{M} be a Gaussian linear mechanism whose output can be used to reconstruct unbiased answers to
 170 the marginals in $Wkload$. For each $\mathbf{A}_i \in Wkload$, let $Var(\mathbf{A}_i; \mathcal{M})$ be the function that returns the
 171 variances of the reconstructed answers to the marginal on \mathbf{A}_i ; the output of $Var(\mathbf{A}_i; \mathcal{M})$ is a vector
 172 v_i with one component for each cell of the marginal on \mathbf{A}_i . A loss function \mathcal{L} aggregates all of these
 173 vectors together: $\mathcal{L}(v_1, \dots, v_k)$. We have the following regularity conditions on the loss function.

174 **DEFINITION 4.1 (Regular Loss Function).** *We say the loss function \mathcal{L} is regular if: (1) \mathcal{L} is convex
 175 and continuous; (2) $\mathcal{L}(v_1, \dots, v_k)$ is minimized when all the v_i are the 0 vectors; and (3) for any i ,
 176 permuting just the components of v_i does not affect the value of $\mathcal{L}(v_1, \dots, v_k)$. This latter condition
 177 just says that cells within the same marginal are equally important.*

178 Loss functions used on prior work are all regular. For example, weighted sum of variances
 179 [32, 52, 30, 53, 37, 51] can be expressed as $\mathcal{L}(v_1, \dots, v_k) = \sum_i c_i \mathbf{1}^T v_i$, where the c_i are the
 180 nonnegative weights that indicate the relative importance of the different marginals. Another pop-
 181 ular loss function is maximum (weighted) variance [46, 18, 42], expressed as $\mathcal{L}(v_1, \dots, v_k) =$
 182 $\max \left\{ \frac{\max(v_1)}{c_1}, \dots, \frac{\max(v_k)}{c_k} \right\}$. Thus, the optimization problem that the selection step needs to solve
 183 is either privacy constrained: minimize loss while keeping privacy cost (defined at the end of Section
 184 2.1) below a threshold γ ; or utility constrained: minimize privacy cost such that the loss is at most γ .

$$\text{Privacy constrained: } \arg \min_{\mathcal{M}} \mathcal{L}(Var(A_1; \mathcal{M}), \dots, Var(A_k; \mathcal{M})) \quad \text{s.t.} \quad pcost(\mathcal{M}) \leq \gamma \quad (1)$$

$$\text{Utility constrained: } \arg \min_{\mathcal{M}} pcost(\mathcal{M}) \quad \text{s.t.} \quad \mathcal{L}(Var(A_1; \mathcal{M}), \dots, Var(A_k; \mathcal{M})) \leq \gamma \quad (2)$$

185 **4.2 Base Mechanisms used by ResidualPlanner**

186 As long as the loss function \mathcal{L} is regular, we will show that an optimal mechanism can be constructed
 187 from a set of base mechanisms that we describe here. We define a *subtraction matrix* \mathbf{Sub}_m to be an
 188 $(m-1) \times m$ matrix where the first column is filled with 1, entries of the form $(i, i+1)$ are -1, and all
 189 other entries are 0. For example, $\mathbf{Sub}_3 = \begin{bmatrix} 1 & -1 & 0 \\ 1 & 0 & -1 \end{bmatrix}$ and $\mathbf{Sub}_2 = \begin{bmatrix} 1 & -1 \end{bmatrix}$. We use these subtraction
 190 matrices to define special matrices called *residual matrices* that are important for our algorithm.

191 For any subset $\mathbf{A} \subseteq \{Att_1, \dots, Att_{n_a}\}$ of attributes, we define the *residual matrix* $\mathbf{R}_{\mathbf{A}}$ as the
 192 Kronecker product $\mathbf{R}_{\mathbf{A}} = \mathbf{V}_1 \otimes \dots \otimes \mathbf{V}_{n_a}$, where $\mathbf{V}_i = \mathbf{1}_{|Att_i|}^T$ if $Att_i \notin \mathbf{A}$ and $\mathbf{V}_i = \mathbf{Sub}_{|Att_i|}$ if
 193 $Att_i \in \mathbf{A}$. Continuing Example 2.1, we have $\mathbf{R}_{\emptyset} = \begin{bmatrix} 1 & 1 \end{bmatrix} \otimes \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}$, and $\mathbf{R}_{\{Att_1\}} = \begin{bmatrix} 1 & -1 \end{bmatrix} \otimes \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}$,
 194 and $\mathbf{R}_{\{Att_2\}} = \begin{bmatrix} 1 & 1 \end{bmatrix} \otimes \begin{bmatrix} 1 & -1 & 0 \\ 1 & 0 & -1 \end{bmatrix}$, and $\mathbf{R}_{\{Att_1, Att_2\}} = \begin{bmatrix} 1 & -1 \end{bmatrix} \otimes \begin{bmatrix} 1 & -1 & 0 \\ 1 & 0 & -1 \end{bmatrix}$.

195 Using subtraction matrices, we also define the matrix $\Sigma_{\mathbf{A}}$ as the Kronecker product
 196 $\bigotimes_{Att_i \in \mathbf{A}} (\mathbf{Sub}_{|Att_i|} \mathbf{Sub}_{|Att_i|}^T)$ and we note that it is proportional to $\mathbf{R}_{\mathbf{A}} \mathbf{R}_{\mathbf{A}}^T$. Σ_{\emptyset} is defined as 1.
 197 Each subset \mathbf{A} of attributes can be associated with a “base” mechanism $\mathcal{M}_{\mathbf{A}}$ that takes as input the
 198 data vector \mathbf{x} and a scalar parameter $\sigma_{\mathbf{A}}^2$ for controlling how noisy the answer is. $\mathcal{M}_{\mathbf{A}}$ is defined as:

$$\mathcal{M}_{\mathbf{A}}(\mathbf{x}; \sigma_{\mathbf{A}}^2) \equiv \mathbf{R}_{\mathbf{A}} \mathbf{x} + N(\mathbf{0}, \sigma_{\mathbf{A}}^2 \Sigma_{\mathbf{A}}) \quad (3)$$

199 The residual matrices $\mathbf{R}_{\mathbf{A}}$ used by base mechanisms form a linearly independent basis that compactly
 200 represent marginals, as the next result shows.

201 **THEOREM 4.2.** *Let \mathbf{A} be a set of attributes and let $\mathbf{Q}_{\mathbf{A}}$ be the matrix representation of the marginal*
 202 *on \mathbf{A} . Then the rows of the matrices $\mathbf{R}_{\mathbf{A}'}$, for all $\mathbf{A}' \subseteq \mathbf{A}$, form a linearly independent basis of the*
 203 *row space of $\mathbf{Q}_{\mathbf{A}}$. Furthermore, if $\mathbf{A}' \neq \mathbf{A}''$ then $\mathbf{R}_{\mathbf{A}'} \mathbf{R}_{\mathbf{A}''}^T = \mathbf{0}$ (they are mutually orthogonal).*

204 **REMARK 4.3.** *To build an intuitive understanding of residual matrices, consider again Example*
 205 *2.1. Both \mathbf{R}_{\emptyset} and \mathbf{Q}_{\emptyset} are the sum query (marginal on no attributes). The rows of $\mathbf{R}_{\{Att_1\}}$ span the*
 206 *subspace of $\mathbf{Q}_{\{Att_1\}}$ that is orthogonal to \mathbf{Q}_{\emptyset} (and similarly for $\mathbf{R}_{\{Att_2\}}$). The rows of $\mathbf{R}_{\{Att_1, Att_2\}}$*
 207 *span the subspace of $\mathbf{Q}_{\{Att_1, Att_2\}}$ that is orthogonal to both $\mathbf{Q}_{\{Att_1\}}$ and $\mathbf{Q}_{\{Att_2\}}$. Hence a residual*
 208 *matrix spans the subspace of a marginal that is orthogonal to its sub-marginals.*

209 Theorem 4.2 has several important implications. If we define the downward closure of a marginal
 210 workload $Wkload = \{\mathbf{A}_1, \dots, \mathbf{A}_k\}$ as the collection of all subsets of the sets in $Wkload$ (i.e.,
 211 $\text{closure}(Wkload) = \{\mathbf{A}' : \mathbf{A}' \subseteq \mathbf{A} \text{ for some } \mathbf{A} \in Wkload\}$) then the theorem implies that the
 212 combined rows from $\{\mathbf{R}_{\mathbf{A}'} : \mathbf{A}' \in \text{closure}(Wkload)\}$ forms a linearly independent basis for the
 213 marginals in the workload. In other words, it is a linearly independent bases for the space spanned by
 214 the rows of the marginal query matrices $\mathbf{Q}_{\mathbf{A}}$ for $\mathbf{A} \in Wkload$. Thus, in order to provide privacy-
 215 preserving answers to all of the marginals represented in $Wkload$, we need all the mechanisms $\mathcal{M}_{\mathbf{A}'}$
 216 for $\mathbf{A}' \in \text{closure}(Wkload)$ – any other matrix mechanism that provides fewer noisy outputs cannot
 217 reconstruct unbiased answers to the workload marginals. This is proved in Theorem 4.4, which also
 218 states that optimality is achieved by carefully setting the $\sigma_{\mathbf{A}}$ noise parameter for each $\mathcal{M}_{\mathbf{A}}$.

219 **THEOREM 4.4.** *Given a marginal workload $Wkload$ and a regular loss function \mathcal{L} , suppose the*
 220 *optimization problem (either Equation 1 or 2) is feasible. Then there exist nonnegative constants*
 221 *$\sigma_{\mathbf{A}}^2$ for each $\mathbf{A} \in \text{closure}(Wkload)$ (the constants do not depend on the data), such that the optimal*
 222 *linear Gaussian mechanism \mathcal{M}_{opt} releases $\mathcal{M}_{\mathbf{A}}(\mathbf{x}; \sigma_{\mathbf{A}}^2)$ for all $\mathbf{A} \in \text{closure}(Wkload)$. Furthermore,*
 223 *any matrix mechanism for this workload must release at least this many noise query answers.*

Algorithm 1: Efficient implementation of $\mathcal{M}_{\mathbf{A}}(\mathbf{x}; \sigma_{\mathbf{A}}^2) \equiv \mathbf{R}_{\mathbf{A}} \mathbf{x} + N(\mathbf{0}, \sigma_{\mathbf{A}}^2 \Sigma_{\mathbf{A}})$

```

1  $\mathbf{v} \leftarrow \mathbf{Q}_{\mathbf{A}} \mathbf{x}$  // Evaluate the true marginal
2  $m \leftarrow \prod_{Att_i \in \mathbf{A}} |Att_i|$ 
3  $\mathbf{H} \leftarrow \bigotimes_{Att_i \in \mathbf{A}} \mathbf{Sub}_{|Att_i|}$  // Use implicit representation, don't expand
4  $\mathbf{z} \leftarrow N(\mathbf{0}, \mathcal{I}_m)$  // independent noise
5 return  $\mathbf{H} \mathbf{v} + \sigma_{\mathbf{A}} \mathbf{H} \mathbf{z}$  // use kron-product/vector multiplication from [37]
```

224 $\mathcal{M}_{\mathbf{A}}$ can be evaluated efficiently, directly from the marginal of \mathbf{x} on attribute set \mathbf{A} , as shown in
 225 Algorithm 1. It uses the technique from [37] to perform fast multiplication between a Kronecker
 226 product and a vector. The privacy cost $\text{pcost}(\mathcal{M}_{\mathbf{A}})$ of each base mechanism $\mathcal{M}_{\mathbf{A}}$ is also easy to
 227 compute and is given by the following theorem.

228 **THEOREM 4.5.** *The privacy cost of $\mathcal{M}_{\mathbf{A}}$ with noise parameter $\sigma_{\mathbf{A}}^2$ is $\frac{1}{\sigma_{\mathbf{A}}^2} \prod_{Att_i \in \mathbf{A}} \frac{|Att_i|-1}{|Att_i|}$ and the
 229 evaluation of $\mathcal{M}_{\mathbf{A}}$ given in Algorithm 1 is correct.*

230 4.3 Reconstruction

231 Next we explain how to reconstruct unbiased answers to marginal queries from the outputs of the base
 232 mechanisms and how to compute (co)variances of the reconstructed marginals efficiently, without any
 233 heavy matrix operations (inversion, pseudo-inverses, etc.). Then, given the closed form expressions
 234 for marginals and privacy cost (Theorem 4.5), we will be able to explain in Section 4.4 how to
 235 optimize the $\sigma_{\mathbf{A}}^2$ parameters of the base mechanisms $\mathcal{M}_{\mathbf{A}}$ to optimize regular loss functions \mathcal{L} .

236 Since the base mechanisms were built using a linearly independent basis, reconstruction is unique –
 237 just efficiently invert the basis. Hence, unlike PGM and its extensions [41, 40] our reconstruction
 238 algorithm does not need to solve an optimization problem and can reconstruct each marginal in-
 239 dependently, thus allowing marginals to be reconstructed in parallel, or as needed by users. The
 240 reconstructed marginals are consistent with each other (any two reconstructed marginals agree on
 241 their sub-marginals). Just as the subtraction matrices \mathbf{Sub}_k were useful in constructing the base
 242 mechanisms $\mathcal{M}_{\mathbf{A}}$, their pseudo-inverses \mathbf{Sub}_k^\dagger are useful for reconstructing noisy marginals from the
 243 noisy answers of $\mathcal{M}_{\mathbf{A}}$. The pseudo-inverses have a closed form. For example $\mathbf{Sub}_4 = \begin{bmatrix} 1 & -1 & 0 & 0 \\ 1 & 0 & -1 & 0 \\ 1 & 0 & 0 & -1 \end{bmatrix}$

244 and $\mathbf{Sub}_4^\dagger = \frac{1}{4} \begin{bmatrix} 1 & 1 & 1 \\ -3 & 1 & 1 \\ 1 & -3 & 1 \\ 1 & 1 & -3 \end{bmatrix}$. More generally, they are expressed as follows:

245 **LEMMA 4.6.** *For any Att_i , let $\ell = |Att_i|$. The matrix \mathbf{Sub}_ℓ has the following block matrix, with
 246 dimensions $\ell \times (\ell - 1)$, as its pseudo-inverse (and right inverse): $\mathbf{Sub}_\ell^\dagger = \frac{1}{\ell} \begin{bmatrix} \mathbf{1}_{\ell-1}^T \\ \mathbf{1}_{\ell-1} \mathbf{1}_{\ell-1}^T - \ell \mathcal{I}_{\ell-1} \end{bmatrix}$.*

247 Each mechanism $\mathcal{M}_{\mathbf{A}}$, for $\mathbf{A} \in \text{closure}(Wkload)$, has a noise scale parameter $\sigma_{\mathbf{A}}^2$ and a noisy output
 248 that we denote by $\omega_{\mathbf{A}}$. After we have obtained the noisy outputs $\omega_{\mathbf{A}}$ for all $\mathbf{A} \in \text{closure}(Wkload)$,
 249 we can proceed with the reconstruction phase. The reconstruction of an unbiased noisy answer for
 250 any marginal on an attribute set $\mathbf{A} \in \text{closure}(Wkload)$ is obtained using Algorithm 2. We note
 251 that to reconstruct a marginal on attribute set \mathbf{A} , one only needs to use the noisy answers $\omega_{\mathbf{A}'}$ for
 252 $\mathbf{A}' \in \text{closure}(\mathbf{A})$. In other words, if we want to reconstruct a marginal on attribute set $\{Att_1, Att_2\}$,
 253 we only need the outputs of \mathcal{M}_\emptyset , $\mathcal{M}_{\{Att_1\}}$, $\mathcal{M}_{\{Att_2\}}$, and $\mathcal{M}_{\{Att_1, Att_2\}}$ no matter how many other
 254 attributes are in the data and no matter what other marginals are in the *Wkload*. We emphasize again,
 255 the reconstruction phase does not run the base mechanisms anymore, it is purely post-processing.

Algorithm 2: Reconstruct Unbiased Answers to the Marginal on \mathbf{A}

Input: Noise scale parameters $\sigma_{\mathbf{A}'}^2$, and noisy answer vector $\omega_{\mathbf{A}'}$ of mechanism $\mathcal{M}_{\mathbf{A}'}$ for
 every $\mathbf{A}' \in \text{closure}(\mathbf{A})$.

Output: \mathbf{q} is output as an unbiased noisy estimate of $\mathbf{Q}_{\mathbf{A}} \mathbf{x}$.

1 $\mathbf{q} \leftarrow \mathbf{0}$

2 **for each** $\mathbf{A}' \in \text{closure}(\mathbf{A})$ **do**

3 $\mathbf{U} \leftarrow \mathbf{V}_1 \otimes \cdots \otimes \mathbf{V}_{n_a}$, where $\mathbf{V}_i = \begin{cases} \mathbf{Sub}_{|Att_i|}^\dagger & \text{if } Att_i \in \mathbf{A}' \\ \frac{1}{|Att_i|} \mathbf{1}_{|Att_i|} & \text{if } Att_i \in \mathbf{A}/\mathbf{A}' \\ [1] & \text{if } Att_i \notin \mathbf{A} \end{cases}$

4 $\mathbf{q} \leftarrow \mathbf{q} + \mathbf{U} \omega_{\mathbf{A}'}$ // use kron-product/vector multiplication from [37]

5 **return** \mathbf{q}

256 **THEOREM 4.7.** *Given a marginal workload $Wkload$ and positive numbers $\sigma_{\mathbf{A}}^2$ for each $\mathbf{A} \in$
 257 $\text{closure}(Wkload)$, let \mathcal{M} be the mechanism that outputs $\{\mathcal{M}_{\mathbf{A}}(\mathbf{x}; \sigma_{\mathbf{A}}^2) : \mathbf{A} \in \text{closure}(Wkload)\}$*

258 and let $\{\omega_{\mathbf{A}} : \mathbf{A} \in \text{closure}(Wkload)\}$ denote the privacy-preserving noisy answers (e.g.,
 259 $\omega_{\mathbf{A}} = \mathcal{M}_{\mathbf{A}}(\mathbf{x}, \sigma^2)$). Then for any marginal on an attribute set $\mathbf{A} \in \text{closure}(Wkload)$, Algo-
 260 rithm 2 returns the unique linear unbiased estimate of $\mathbf{Q}_{\mathbf{A}}\mathbf{x}$ (i.e., answers to the marginal query)
 261 that can be computed from the noisy differentially private answers.

262 The variances $\text{Var}(\mathbf{A}; \mathcal{M})$ of all the noisy cell counts of the marginal on \mathbf{A} is the vector
 263 whose components are all equal to $\sum_{\mathbf{A}' \subseteq \mathbf{A}} \left(\sigma_{\mathbf{A}'}^2 \prod_{\text{Att}_i \in \mathbf{A}'} \frac{|\text{Att}_i| - 1}{|\text{Att}_i|} * \prod_{\text{Att}_j \in (\mathbf{A}/\mathbf{A}')} \frac{1}{|\text{Att}_j|^2} \right)$.
 264 The covariance between any two noisy answers of the marginal on \mathbf{A} is
 265 $\sum_{\mathbf{A}' \subseteq \mathbf{A}} \left(\sigma_{\mathbf{A}'}^2 \prod_{\text{Att}_i \in \mathbf{A}'} \frac{-1}{|\text{Att}_i|} * \prod_{\text{Att}_j \in (\mathbf{A}/\mathbf{A}')} \frac{1}{|\text{Att}_j|^2} \right)$.

266 4.4 Optimizing the Base Mechanism Selection

267 We now consider how to find the optimal Gaussian linear mechanism \mathcal{M}^* that solves the optimization
 268 problems in Equations 1 or 2. Given a workload on marginals $Wkload$, the optimization involves
 269 $\text{Var}(\mathbf{A}; \mathcal{M}^*)$ for $\mathbf{A} \in Wkload$ (the variance of the marginal answers reconstructed from the output
 270 of \mathcal{M}^*) and $\text{pcost}(\mathcal{M}^*)$, from which the privacy parameters of different flavors of differential privacy
 271 can be computed.

272 Theorem 4.4 says that \mathcal{M}^* works by releasing $\mathcal{M}_{\mathbf{A}}(\mathbf{x}; \sigma_{\mathbf{A}}^2)$ for each $\mathbf{A} \in \text{closure}(Wkload)$ for
 273 appropriately chosen values of $\sigma_{\mathbf{A}}^2$. The privacy cost $\text{pcost}(\mathcal{M}^*)$ is the sum of the privacy costs of
 274 the $\mathcal{M}_{\mathbf{A}}$. Theorem 4.5 therefore shows that $\text{pcost}(\mathcal{M}^*)$ is a positive linear combination of the values
 275 $1/\sigma_{\mathbf{A}}^2$ for $\mathbf{A} \in \text{closure}(Wkload)$ and is therefore convex in the $\sigma_{\mathbf{A}}^2$ values. Meanwhile, Theorem 4.7
 276 shows how to represent, for each $\mathbf{A} \in \text{closure}(Wkload)$, the quantity $\text{Var}(\mathbf{A}; \mathcal{M}^*)$ as a positive
 277 linear combination of $\sigma_{\mathbf{A}'}^2$, for $\mathbf{A}' \in \text{closure}(\mathbf{A}) \subseteq \text{closure}(Wkload)$. Therefore, the loss function \mathcal{L}
 278 is also convex in the $\sigma_{\mathbf{A}}^2$ values.

279 Thus the optimization problems in Equations 1 and 2 can be written as minimizing a convex function
 280 of the $\sigma_{\mathbf{A}}^2$ subject to convex constraints. In fact, in Equation 2, the constraints are linear when
 281 the optimization variables represent the $\sigma_{\mathbf{A}}^2$ and in Equation 1 the constraints are linear when the
 282 optimization variables represent the $1/\sigma_{\mathbf{A}}^2$. Furthermore, when the loss function is the weighted sum
 283 of variances of the marginal cells, the solution can be obtained in closed form (see supplementary
 284 material). Otherwise, we use CVXPY/ECOS [12, 14] for solving these convex optimization problems.

285 4.5 Computational Complexity

286 The time complexity of the steps of our framework is provided in the following theorem. It can be
 287 expressed in terms of the sizes of the marginals the user is asking for. Crucially, it does *not* depend
 288 on the universe size $|\text{Att}_1| \times \dots \times |\text{Att}_{n_a}|$, which accounts for the scalability.

289 THEOREM 4.8. Let n_a be the total number of attributes. Let $\#\text{cells}(\mathbf{A})$ denote the number of cells in
 290 the marginal on attribute set \mathbf{A} . Then:

- 291 1. Expressing the privacy cost of the optimal mechanism \mathcal{M}^* as a linear combination of the $1/\sigma_{\mathbf{A}}^2$
 292 values takes $O(\sum_{\mathbf{A} \in Wkload} \#\text{cells}(\mathbf{A}))$ total time.
- 293 2. Expressing all of the $\text{Var}(\mathbf{A}; \mathcal{M}^*)$, for $\mathbf{A} \in Wkload$, as a linear combinations of the $\sigma_{\mathbf{A}}^2$ values
 294 can be done in $O(\sum_{\mathbf{A} \in Wkload} \#\text{cells}(\mathbf{A}))$ total time.
- 295 3. Computing all the noisy outputs of the optimal mechanism (i.e., $\mathcal{M}_{\mathbf{A}}(\mathbf{x}; \sigma_{\mathbf{A}}^2)$ for $\mathbf{A} \in$
 296 $\text{closure}(Wkload)$) takes $O(n_a \sum_{\mathbf{A} \in Wkload} \prod_{\text{Att}_i \in \mathbf{A}} (|\text{Att}_i| + 1))$ total time after the true an-
 297 swers have been precomputed (Line 1 in Algorithm 1). Note that the total number of cells on
 298 marginals in $Wkload$ is $O(\sum_{\mathbf{A} \in Wkload} \prod_{\text{Att}_i \in \mathbf{A}} |\text{Att}_i|)$.
- 299 4. Reconstructing marginals for all $\mathbf{A} \in Wkload$ takes $O(\sum_{\mathbf{A} \in Wkload} |\mathbf{A}| \#\text{cells}(\mathbf{A})^2)$ total time.
- 300 5. Computing the variance of the cells for all of the marginals for $\mathbf{A} \in Wkload$ can be done in
 301 $O(\sum_{\mathbf{A} \in Wkload} \#\text{cells}(\mathbf{A}))$ total time.

302 To get a sense of these numbers, consider a dataset with 20 attributes, each having 3 possible values.
 303 If the workload consists of all 3-way marginals, there are 1,140 marginals each with 27 cells so
 304 $n_{\text{cells}} = 30,780$. The quantity inside the big-O for the selection step is 1,459,200 (roughly the
 305 number of scalar multiplications needed). These are all easily manageable on modern computers
 306 even without GPUs. Our experiments, under more challenging conditions, run in seconds.

307 **5 Experiments**

308 We next compare the accuracy and scalability of ResidualPlanner against HDMM [38], including
 309 variations of HDMM with faster reconstruction phases [41]. The hardware used was an Ubuntu
 310 22.04.2 server with 12 Intel(R) Core(TM) i7-8700 CPU @ 3.20GHz processors and 32GB of
 311 DDR4 RAM. We use 3 real datasets to evaluate accuracy and 1 synthetic dataset to evaluate scal-
 312 ability. The real datasets are (1) the Adult dataset [16] with 14 attributes, each having domain
 313 sizes 100, 100, 100, 99, 85, 42, 16, 15, 9, 7, 6, 5, 2, 2, respectively, resulting in a record domain size
 314 of $6.41 * 10^{17}$; (2) the CPS dataset [9] with 5 attributes, each having domain size 100, 50, 7, 4, 2,
 315 respectively, resulting in a record domain size of $2.8 * 10^5$; (3) the Loans dataset [24] with 12
 316 attributes, each having domain size 101, 101, 101, 101, 3, 8, 36, 6, 51, 4, 5, 15, respectively, resulting
 317 in a record domain size of $8.25 * 10^{15}$. The synthetic dataset is called Synth- n^d . Here d refers to
 318 the number of attributes (we experiment from $d = 2$ to $d = 100$) and n is the domain size of each
 319 attribute. The running times of the algorithms only depend on n and d and not on the records in
 320 the synthetic data. For all experiments, we set the privacy cost $pcost$ to 1, so all mechanisms being
 321 compared satisfy 0.5-zCDP and 1-Gaussian DP.

322 **5.1 Scalability of the Selection Phase**

323 We first consider how long each method takes to perform the selection phase (i.e., determine what
 324 needs noisy answers and how much noise to use). HDMM can only optimize total variance, which
 325 is equivalent to root mean squared error. For ResidualPlanner we consider both RMSE and max
 326 variance as objectives (the latter is a harder to solve problem). Each algorithm is run 5 times and the
 327 average is taken. Table 1 shows running time results; accuracy results will be presented later.

Table 1: **Time for Selection Step in seconds** on Synth- n^d dataset. $n = 10$ and the number of attributes d varies. The workload consists of all marginals on ≤ 3 attributes each. Times for HDMM are reported with ± 2 standard deviations.

d	HDMM RMSE Objective	ResidualPlanner RMSE Objective	ResidualPlanner Max Variance Objective
2	0.013 \pm 0.003	0.001 \pm 0.0008	0.007 \pm 0.001
6	0.065 \pm 0.012	0.002 \pm 0.0008	0.009 \pm 0.001
10	0.639 \pm 0.059	0.009 \pm 0.001	0.018 \pm 0.001
12	4.702 \pm 0.315	0.015 \pm 0.001	0.028 \pm 0.001
14	46.054 \pm 12.735	0.025 \pm 0.002	0.041 \pm 0.001
15	201.485 \pm 13.697	0.030 \pm 0.017	0.050 \pm 0.001
20	Out of memory	0.079 \pm 0.017	0.123 \pm 0.023
30	Out of memory	0.247 \pm 0.019	0.461 \pm 0.024
50	Out of memory	1.207 \pm 0.047	4.011 \pm 0.112
100	Out of memory	9.913 \pm 0.246	121.224 \pm 3.008

328 As we can see, optimizing for max variance is more difficult than for RMSE, but ResidualPlanner
 329 does it quickly even for data settings too big for HDMM. The runtime of HDMM increases rapidly,
 330 while even for the extreme end of our experiments, ResidualPlanner needs just a few minutes.

331 **5.2 Scalability of the Reconstruction Phase**

332 We next evaluate the scalability of the reconstruction phase under the same settings. The reconstruc-
 333 tion speed for ResidualPlanner does not depend on the objective of the selection phase. Here we
 334 compare against HDMM [38] and a version of HDMM with improved reconstruction scalability
 335 called HDMM+PGM [38, 41] (the PGM settings used 50 iterations of its Local-Inference estimator,
 336 as the default 1000 was too slow). Since HDMM cannot perform the selection phase after a certain
 337 point, reconstruction results also become unavailable. Table 2 shows ResidualPlanner is clearly faster.

338 **5.3 Accuracy Comparisons**

339 Since ResidualPlanner is optimal, the purpose of the accuracy comparisons is a sanity check. For
 340 RMSE, we compare the quality of ResidualPlanner to the theoretically optimal lower bound known

Table 2: **Time for Reconstruction Step in seconds** on Synth- n^d dataset. $n = 10$ and the number of attributes d varies. The workload consists of all marginals on ≤ 3 attributes each. Times are reported with ± 2 standard deviations. Reconstruction can only be performed if the select step completed.

d	HDMM	HDMM + PGM	ResidualPlanner
2	0.003 ± 0.0006	0.155 ± 0.011	0.005 ± 0.003
6	0.173 ± 0.011	4.088 ± 0.233	0.023 ± 0.004
10	Out of memory in reconstruction	20.340 ± 2.264	0.125 ± 0.032
12	Out of memory in reconstruction	39.162 ± 1.739	0.207 ± 0.004
14	Out of memory in reconstruction	69.975 ± 4.037	0.330 ± 0.006
15	Out of memory in reconstruction	91.101 ± 7.621	0.413 ± 0.006
20	Unavailable (select step failed)	Unavailable (select step failed)	1.021 ± 0.011
30	Unavailable (select step failed)	Unavailable (select step failed)	3.587 ± 0.053
50	Unavailable (select step failed)	Unavailable (select step failed)	17.029 ± 0.212
100	Unavailable (select step failed)	Unavailable (select step failed)	154.538 ± 15.045

341 as the SVD bound [31] (they match, as shown in Table 3). We note ResidualPlanner can provide
 342 solutions even when the SVD bound is infeasible to compute. Then we compare ResidualPlanner
 343 to HDMM when the user is interested in the maximum variance objective. This just shows that it
 344 is important to optimize for the user’s objective function and that the optimal solution for RMSE
 345 (the only objective HDMM can optimize) is not a good general-purpose approximation for other
 346 objectives (as shown in Table 4). Additional comparisons are provided in the supplementary material.

Table 3: RMSE Comparisons to the theoretical lower bound SVD Bound [31]

Workload	Adult Dataset		CPS Dataset		Loans Dataset	
	ResPlan	SVDB	ResPlan	SVDB	ResPlan	SVDB
1-way Marginals	3.047	3.047	1.744	1.744	2.875	2.875
2-way Marginals	6.359	6.359	2.035	2.035	5.634	5.634
3-way Marginals	10.515	10.515	2.048	2.048	8.702	8.702
≤ 3 -way Marginals	10.665	10.665	2.276	2.276	8.876	8.876

Table 4: Max Variance Comparisons with ResidualPlanner and HDMM (showing that being restricted to optimizing only RMSE is not a good approximation of Max Variance optimization).

Workload	Adult Dataset		CPS Dataset		Loans Dataset	
	ResPlan	HDMM	ResPlan	HDMM	ResPlan	HDMM
1-way Marginals	12.047	41.772	4.346	13.672	10.640	33.256
2-way Marginals	67.802	599.843	7.897	47.741	52.217	437.478
3-way Marginals	236.843	5675.238	7.706	71.549	156.638	3095.997
≤ 3 -way Marginals	253.605	6677.253	13.216	415.073	180.817	4317.709

347 6 Limitations, Conclusion, and Future Work.

348 In this paper, we introduced ResidualPlanner, a matrix mechanism that is scalable and optimal for
 349 marginals under Gaussian noise, for a large class of convex objective functions. While these are
 350 important improvements to the state of the art, there are limitations.

351 First, for some attributes, a user might not want marginals. For example, they might want range
 352 queries or queries with hierarchies (e.g., how many people drive sedans vs. vans; out of the sedans,
 353 how many are red vs. green, etc) [2, 28, 36]. In some cases, an attribute might have an infinite domain
 354 (e.g., a URL) and need to be handled differently [27, 45]. In other cases, the user may want other
 355 noise distributions, like the Laplace. These types of queries do not have the same type of symmetry
 356 as marginals that was crucial to proving the optimality of ResidualPlanner. For these situations, one
 357 of the key ideas of ResidualPlanner can be used – find a linear basis that compactly represents both
 358 the queries and “residual” (information provided by a query that is not contained in the other queries).
 359 Such a feature would result in scalability. It is future work to determine how to extend both scalability
 360 and optimality to such situations.

References

- 361
- 362 [1] Nazmiye Ceren Abay, Yan Zhou, Murat Kantarcioglu, Bhavani Thuraisingham, and Latanya
363 Sweeney. Privacy preserving synthetic data release using deep learning. In *Machine Learning
364 and Knowledge Discovery in Databases: European Conference, ECML PKDD 2018, Dublin,
365 Ireland, September 10–14, 2018, Proceedings, Part I 18*, pages 510–526. Springer, 2019.
- 366 [2] John M. Abowd, Robert Ashmead, Ryan Cumings-Menon, Simson Garfinkel, Micah Hei-
367 neck, Christine Heiss, Robert Johns, Daniel Kifer, Philip Leclerc, Ashwin Machanava-
368 jjhala, Brett Moran, William Sexton, Matthew Spence, and Pavel Zhuravlev. The 2020
369 census disclosure avoidance system topdown algorithm. *Harvard Data Science Re-
370 view*, forthcoming. Preprint [https://www.census.gov/library/working-papers/2022/
371 adrm/CED-WP-2022-002.html](https://www.census.gov/library/working-papers/2022/adrm/CED-WP-2022-002.html).
- 372 [3] Hassan Jameel Asghar, Ming Ding, Thierry Rakotoarivelo, Sirine Mrabet, and Mohamed Ali
373 Kaafar. Differentially private release of high-dimensional datasets using the gaussian copula.
374 *arXiv preprint arXiv:1902.01499*, 2019.
- 375 [4] Sergul Aydore, William Brown, Michael Kearns, Krishnam Kenthapadi, Luca Melis, Aaron
376 Roth, and Ankit A Siva. Differentially private query release through adaptive projection. In
377 *International Conference on Machine Learning*, pages 457–467. PMLR, 2021.
- 378 [5] Borja Balle and Yu-Xiang Wang. Improving the gaussian mechanism for differential privacy:
379 Analytical calibration and optimal denoising. In *International Conference on Machine Learning,
380 ICML*, 2018.
- 381 [6] Boaz Barak, Kamalika Chaudhuri, Cynthia Dwork, Satyen Kale, Frank McSherry, and Kunal
382 Talwar. Privacy, accuracy, and consistency too: a holistic solution to contingency table release.
383 In *Proceedings of the twenty-sixth ACM SIGMOD-SIGACT-SIGART symposium on Principles
384 of database systems*, pages 273–282, 2007.
- 385 [7] Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions,
386 and lower bounds. In *Proceedings, Part I, of the 14th International Conference on Theory of
387 Cryptography - Volume 9985*, 2016.
- 388 [8] U.S. Census Bureau. Decennial census: 2010 summary files. [https://www.census.gov/
389 mp/www/cat/decennial_census_2010/](https://www.census.gov/mp/www/cat/decennial_census_2010/).
- 390 [9] U.S. Census Bureau. The current population survey (cps). [https://www.census.gov/
391 programs-surveys/cps.html](https://www.census.gov/programs-surveys/cps.html), 2023.
- 392 [10] Kuntai Cai, Xiaoyu Lei, Jianxin Wei, and Xiaokui Xiao. Data synthesis via differentially private
393 markov random fields. *Proceedings of the VLDB Endowment*, 14(11):2190–2202, 2021.
- 394 [11] Rui Chen, Qian Xiao, Yu Zhang, and Jianliang Xu. Differentially private high-dimensional
395 data publication via sampling-based inference. In *Proceedings of the 21th ACM SIGKDD
396 international conference on knowledge discovery and data mining*, pages 129–138, 2015.
- 397 [12] Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for
398 convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.
- 399 [13] Bolin Ding, Marianne Winslett, Jiawei Han, and Zhenhui Li. Differentially private data
400 cubes: Optimizing noise sources and consistency. In *Proceedings of the 2011 ACM SIGMOD
401 International Conference on Management of Data*, 2011.
- 402 [14] Alexander Domahidi, Eric Chu, and Stephen Boyd. Ecos: An socp solver for embedded systems.
403 In *2013 European control conference (ECC)*, pages 3071–3076. IEEE, 2013.
- 404 [15] Jinshuo Dong, Aaron Roth, and Weijie J. Su. Gaussian differential privacy. *Journal of the Royal
405 Statistical Society: Series B (Statistical Methodology)*, 84(1):3–37, 2022.
- 406 [16] Dheeru Dua and Casey Graff. UCI machine learning repository, 2019.

- 407 [17] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our
408 data, ourselves: Privacy via distributed noise generation. In *EUROCRYPT*, pages 486–503,
409 2006.
- 410 [18] Alexander Edmonds, Aleksandar Nikolov, and Jonathan Ullman. The power of factorization
411 mechanisms in local and central differential privacy. In *Proceedings of the 52nd Annual ACM*
412 *SIGACT Symposium on Theory of Computing*, pages 425–438, 2020.
- 413 [19] Marco Gaboardi, Emilio Jesús Gallego Arias, Justin Hsu, Aaron Roth, and Zhiwei Steven
414 Wu. Dual query: Practical private query release for high dimensional data. In *International*
415 *Conference on Machine Learning*, pages 1170–1178. PMLR, 2014.
- 416 [20] Marco Gaboardi, Hyun Lim, Ryan Rogers, and Salil Vadhan. Differentially private chi-squared
417 hypothesis testing: Goodness of fit and independence testing. In *International conference on*
418 *machine learning*, pages 2111–2120. PMLR, 2016.
- 419 [21] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2023.
- 420 [22] Moritz Hardt, Katrina Ligett, and Frank McSherry. A simple and practical algorithm for
421 differentially private data release. *Advances in neural information processing systems*, 25, 2012.
- 422 [23] James Jordon, Jinsung Yoon, and Mihaela Van Der Schaar. Pate-gan: Generating synthetic data
423 with differential privacy guarantees. In *International conference on learning representations*,
424 2019.
- 425 [24] Kaggle. Loan prediction problem dataset. [https://www.kaggle.com/
426 altruistdelhite04/loan-prediction-problem-dataset](https://www.kaggle.com/altruistdelhite04/loan-prediction-problem-dataset), 2021. Accessed: May
427 8th, 2023.
- 428 [25] Kazuya Kakizaki, Kazuto Fukuchi, and Jun Sakuma. Differentially private chi-squared test by
429 unit circle mechanism. In *International Conference on Machine Learning*, pages 1761–1770.
430 PMLR, 2017.
- 431 [26] Daniel Kifer and Ryan Rogers. A new class of private chi-square tests. In *Proceedings of the*
432 *20th International Conference on Artificial Intelligence and Statistics, AISTATS*, volume 17,
433 pages 991–1000, 2016.
- 434 [27] Aleksandra Korolova, Krishnaram Kenthapadi, Nina Mishra, and Alexandros Ntoulas. Releasing
435 search queries and clicks privately. In *Proceedings of the 18th international conference on*
436 *World wide web*, pages 171–180, 2009.
- 437 [28] Yu-Hsuan Kuo, Cho-Chun Chiu, Daniel Kifer, Michael Hay, and Ashwin Machanavajjhala.
438 Differentially private hierarchical count-of-counts histograms. *arXiv preprint arXiv:1804.00370*,
439 2018.
- 440 [29] Jing Lei. Differentially private m-estimators. *Advances in Neural Information Processing*
441 *Systems*, 24, 2011.
- 442 [30] Chao Li and Gerome Miklau. An adaptive mechanism for accurate query answering under
443 differential privacy. *Proceedings of the VLDB Endowment*, 5(6), 2012.
- 444 [31] Chao Li and Gerome Miklau. Optimal error of query sets under the differentially-private matrix
445 mechanism. In *Proceedings of the 16th International Conference on Database Theory*, pages
446 272–283, 2013.
- 447 [32] Chao Li, Gerome Miklau, Michael Hay, Andrew McGregor, and Vibhor Rastogi. The matrix
448 mechanism: Optimizing linear counting queries under differential privacy. *The VLDB Journal*,
449 24(6):757–781, December 2015.
- 450 [33] Haoran Li, Li Xiong, and Xiaoqian Jiang. Differentially private synthesization of multi-
451 dimensional data using copula functions. In *Advances in database technology: proceedings.*
452 *International conference on extending database technology*, volume 2014, page 475. NIH Public
453 Access, 2014.

- 454 [34] Terrance Liu, Giuseppe Vietri, Thomas Steinke, Jonathan Ullman, and Steven Wu. Leveraging
455 public data for practical private query release. In *International Conference on Machine Learning*,
456 pages 6968–6977. PMLR, 2021.
- 457 [35] Terrance Liu, Giuseppe Vietri, and Steven Z Wu. Iterative methods for private synthetic data:
458 Unifying framework and new methods. *Advances in Neural Information Processing Systems*,
459 34, 2021.
- 460 [36] Terrance Liu and Zhiwei Steven Wu. Private synthetic data with hierarchical structure. *arXiv*
461 *preprint arXiv:2206.05942*, 2022.
- 462 [37] Ryan McKenna, Gerome Miklau, Michael Hay, and Ashwin Machanavajjhala. Optimizing error
463 of high-dimensional statistical queries under differential privacy. *Proceedings of the VLDB*
464 *Endowment*, 11(10), 2018.
- 465 [38] Ryan McKenna, Gerome Miklau, Michael Hay, and Ashwin Machanavajjhala. Hdmm: Opti-
466 mizing error of high-dimensional statistical queries under differential privacy. *arXiv preprint*
467 *arXiv:2106.12118*, 2021.
- 468 [39] Ryan McKenna, Brett Mullins, Daniel Sheldon, and Gerome Miklau. Aim: An adaptive and
469 iterative mechanism for differentially private synthetic data. *arXiv preprint arXiv:2201.12677*,
470 2022.
- 471 [40] Ryan McKenna, Siddhant Pradhan, Daniel R Sheldon, and Gerome Miklau. Relaxed marginal
472 consistency for differentially private query answering. *Advances in Neural Information Pro-*
473 *cessing Systems*, 34:20696–20707, 2021.
- 474 [41] Ryan McKenna, Daniel Sheldon, and Gerome Miklau. Graphical-model based estimation and
475 inference for differential privacy. In *International Conference on Machine Learning*, pages
476 4435–4444. PMLR, 2019.
- 477 [42] Aleksandar Nikolov. *New computational aspects of discrepancy theory*. PhD thesis, Rutgers
478 University-Graduate School-New Brunswick, 2014.
- 479 [43] Wahbeh Qardaji, Weining Yang, and Ninghui Li. Priview: practical differentially private release
480 of marginal contingency tables. In *Proceedings of the 2014 ACM SIGMOD international*
481 *conference on Management of data*, pages 1435–1446, 2014.
- 482 [44] Giuseppe Vietri, Grace Tian, Mark Bun, Thomas Steinke, and Steven Wu. New oracle-efficient
483 algorithms for private synthetic data release. In *International Conference on Machine Learning*,
484 pages 9765–9774. PMLR, 2020.
- 485 [45] Royce Wilson, Celia Yuxin Zhang, William Lam, Damien Desfontaines, Daniel Simmons-
486 Marengo, and Bryant Gipson. Differentially private sql with bounded user contribution. In
487 *Proceedings on Privacy Enhancing Technologies Symposium*, 2020.
- 488 [46] Yingtai Xiao, Zeyu Ding, Yuxin Wang, Danfeng Zhang, and Daniel Kifer. Optimizing fitness-
489 for-use of differentially private linear queries. In *VLDB*, 2021.
- 490 [47] Yingtai Xiao, Guanhong Wang, Danfeng Zhang, and Daniel Kifer. Answering private linear
491 queries adaptively using the common mechanism. <https://arxiv.org/abs/2212.00135>,
492 2022.
- 493 [48] Chugui Xu, Ju Ren, Yaoxue Zhang, Zhan Qin, and Kui Ren. Dppro: Differentially private high-
494 dimensional data release via random projection. *IEEE Transactions on Information Forensics*
495 *and Security*, 12(12):3081–3093, 2017.
- 496 [49] Grigory Yaroslavtsev, Graham Cormode, Cecilia M Procopiuc, and Divesh Srivastava. Ac-
497 curate and efficient private release of datacubes and contingency tables. In *2013 IEEE 29th*
498 *International Conference on Data Engineering (ICDE)*, pages 745–756. IEEE, 2013.
- 499 [50] Fei Yu, Stephen E Fienberg, Aleksandra B Slavković, and Caroline Uhler. Scalable privacy-
500 preserving data sharing methodology for genome-wide association studies. *Journal of biomedical*
501 *informatics*, 50:133–141, 2014.

- 502 [51] Ganzhao Yuan, Yin Yang, Zhenjie Zhang, and Zhifeng Hao. Convex optimization for linear
503 query processing under approximate differential privacy. In *Proceedings of the 22nd ACM*
504 *SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016.
- 505 [52] Ganzhao Yuan, Zhenjie Zhang, Marianne Winslett, Xiaokui Xiao, Yin Yang, and Zhifeng Hao.
506 Low-rank mechanism: Optimizing batch queries under differential privacy. *Proc. VLDB Endow.*,
507 5(11):1352–1363, July 2012.
- 508 [53] Ganzhao Yuan, Zhenjie Zhang, Marianne Winslett, Xiaokui Xiao, Yin Yang, and Zhifeng
509 Hao. Optimizing batch linear queries under exact and approximate differential privacy. *ACM*
510 *Transactions on Database Systems (TODS)*, 40(2):1–47, 2015.
- 511 [54] Jun Zhang, Graham Cormode, Cecilia M. Procopiuc, Divesh Srivastava, and Xiaokui Xiao.
512 Privbayes: Private data release via bayesian networks. *ACM Trans. Database Syst.*, 42(4), oct
513 2017.
- 514 [55] Wei Zhang, Jingwen Zhao, Fengqiong Wei, and Yunfang Chen. Differentially private high-
515 dimensional data publication via markov network. *EAI Endorsed Transactions on Security and*
516 *Safety*, 6(19), 2019.
- 517 [56] Zhikun Zhang, Tianhao Wang, Jean Honorio, Ninghui Li, Michael Backes, Shibo He, Jiming
518 Chen, and Yang Zhang. Privsyn: Differentially private data synthesis. 2021.

519	Contents	
520	1 Introduction	1
521	2 Preliminaries	2
522	2.1 Differential Privacy	2
523	2.2 Matrix Mechanism	3
524	3 Additional Related Work	3
525	4 ResidualPlanner	4
526	4.1 Loss Functions Supported by ResidualPlanner	4
527	4.2 Base Mechanisms used by ResidualPlanner	5
528	4.3 Reconstruction	6
529	4.4 Optimizing the Base Mechanism Selection	7
530	4.5 Computational Complexity	7
531	5 Experiments	8
532	5.1 Scalability of the Selection Phase	8
533	5.2 Scalability of the Reconstruction Phase	8
534	5.3 Accuracy Comparisons	8
535	6 Limitations, Conclusion, and Future Work.	9
536	A Table of Notation	15
537	B A Run-through of Residual Planner	16
538	B.1 A Small Dataset and its Vectorized Representation	16
539	B.2 The Marginal Workload and its Representation as a Query Matrix.	16
540	B.3 The Base Mechanisms	18
541	B.4 Reconstruction	19
542	B.5 Privacy Cost and Marginal Variances	20
543	B.6 The Sum-of-Variiances Loss Function	21
544	C Optimality Proof of ResidualPlanner	22
545	C.1 Notation Review	22
546	C.2 Permutations	23
547	C.3 From permutations to interpretations	25
548	D The other proofs about base mechanisms	31
549	E Proofs related to the reconstruction step	33
550	F Computational Complexity Proofs	36

551	G Closed Form Solution to the Weighted Sum of Variances Loss	38
552	H Additional Experiments	38
553	H.1 Scalability	39
554	H.2 Comparison on Real Datasets.	42
555	A Table of Notation	

Table 5: Table of Notation

\mathcal{D} :	Dataset
r_i :	i^{th} record in \mathcal{D}
n_a :	number of attributes each record has
Att_j :	j^{th} attribute.
$ Att_j $:	size of the domain of attribute Att_j .
$a_1^{(j)}, \dots, a_{ Att_j }^{(j)}$:	possible values (domain) of Att_j .
d :	Number of possible records: $d = \prod_{j=1}^{n_a} Att_j $
\mathbf{x} :	Representation of \mathcal{D} as a d -dimensional vector of counts (e.g., histogram)
\mathbf{A} :	(Sub)set of attributes
$\mathbf{Q}_{\mathbf{A}}$:	Matrix representation of the marginal on \mathbf{A} . The counts in the marginal are the result of matrix-vector product $\mathbf{Q}_{\mathbf{A}}\mathbf{x}$.
$\#cells(\mathbf{A})$:	Number of cells in the marginal on \mathbf{A} . Equals $\prod_{Att_i \in \mathbf{A}} Att_i $
e_i :	one-hot encoding vector with entry i being 1 and the rest 0
$e_{i,j}$:	equal to $e_i - e_j$
$\mathbf{1}_k$:	the k -dimensional vector whose entries are all 1.
\mathcal{I}_k :	the $k \times k$ identity matrix
\mathcal{M} :	A privacy mechanism.
ω :	Output of a mechanism.
\mathbf{B} :	Query matrix of a Gaussian linear query mechanism: $\mathcal{M}(\mathbf{x}) \equiv \mathbf{B}\mathbf{x} + N(\mathbf{0}, \Sigma)$
Σ :	Covariance matrix.
$pcost(\mathcal{M})$:	Privacy cost of a Gaussian linear mechanism $\mathcal{M}(\mathbf{x}) \equiv \mathbf{B}\mathbf{x} + N(\mathbf{0}, \Sigma)$. It is defined as the largest diagonal of $\mathbf{B}^T \Sigma^{-1} \mathbf{B}$. Differential privacy parameters can be computed from $pcost(\mathcal{M})$.
$Wkload$:	A workload of marginals. Each element of $Wkload$ is a set of attributes (representing the marginal on those attributes).
n_{cells} :	Total number of cells in the marginals in the marginal workload (i.e., the output size).
$\text{closure}(Wkload)$:	The set of all subsets of $Wkload$. Formally defined as $\{\mathbf{A}' : \mathbf{A}' \subseteq \mathbf{A} \text{ for some } \mathbf{A} \in Wkload\}$.
$Var(\mathbf{A}, \mathcal{M})$:	When the output of \mathcal{M} is used to reconstruct answers to the marginal on \mathbf{A} , then Var returns the vector of variances of the marginal cells.
\mathcal{L} :	The loss function
\dagger :	The operator that gives the pseudo-inverse of a matrix
Sub_m :	An $(m-1) \times m$ subtraction matrix. The first column is filled with 1, entries of the form $(i, i+1)$ are -1, and all other entries are 0.
$\mathbf{R}_{\mathbf{A}}$:	Residual matrix. Given a set $\mathbf{A} \subset \{Att_1, \dots, Att_{n_a}\}$ of attributes, $\mathbf{R}_{\mathbf{A}} = \mathbf{V}_1 \otimes \dots \otimes \mathbf{V}_{n_a}$, where $\mathbf{V}_i = \mathbf{1}_{ Att_i }$ if $Att_i \notin \mathbf{A}$ and $\mathbf{V}_i = \text{Sub}_{ Att_i }$ if $Att_i \in \mathbf{A}$.
$\Sigma_{\mathbf{A}}$:	The covariance matrix used by the base mechanisms, formed as the kronecker product $\bigotimes_{Att_i \in \mathbf{A}} (\text{Sub}_{ Att_i } \text{Sub}_{ Att_i }^T)$. Also $\Sigma_{\emptyset} = 1$.
$\sigma_{\mathbf{A}}$:	Data-independent noise scale parameter
$\mathcal{M}_{\mathbf{A}}$:	The base mechanism defined as $\mathcal{M}_{\mathbf{A}}(\mathbf{x}) \equiv \mathbf{R}_{\mathbf{A}}\mathbf{x} + N(\mathbf{0}, \sigma_{\mathbf{A}}^2 \Sigma_{\mathbf{A}})$. It uses a data-independent noise parameter $\sigma_{\mathbf{A}}^2$
$\omega_{\mathbf{A}}$:	noisy output of mechanism $\mathcal{M}_{\mathbf{A}}$

556 B A Run-through of Residual Planner

557 In this section, we provide a complete runthrough of ResidualPlanner using a small toy dataset.

558 B.1 A Small Dataset and its Vectorized Representation

559 In our example, we have a dataset with 3 attributes, so $n_a = 3$. Att_1 takes values ‘a’ or ‘b’; Att_2
560 takes values ‘y’ or ‘n’; Att_3 takes values 1 or 2 or 3.

In this dataset, there are 5 people, and the tabular representation is shown in Table 6. For each

Att_1	Att_2	Att_3
a	n	2
b	n	3
b	y	3
a	n	2
b	y	3

Table 6: A Toy Dataset \mathcal{D}

561 attribute, we can one-hot encode its attribute values as row vectors. So, for Att_1 , the attribute value
562 ‘a’ is encoded as $[1, 0]$ and ‘b’ is encoded as $[0, 1]$. For Att_2 , the attribute value ‘y’ is encoded as
563 $[1, 0]$ and ‘n’ is encoded as $[0, 1]$. For attribute Att_3 , the attribute value ‘1’ is encoded as $[1, 0, 0]$, the
564 value ‘2’ is encoded as $[0, 1, 0]$ and ‘3’ is encoded as $[0, 0, 1]$.
565

566 The kronecker product representation of a record is the kronecker product of the one-hot encoding
567 of each attributes. So, for example, the record ‘an2’ is encoded as the kronecker product
568 $[1, 0] \otimes [0, 1] \otimes [0, 1, 0]$. When this kronecker product is expanded, it has 12 components. One of the
569 contains a 1 and the rest contain a 0. Thus the expanded kronecker product can be thought of as a
570 one-hot encoding of the entire record.

571 Indeed, in the expanded kronecker product, each dimension of the resulting vector is associated with
572 a record. In table 7, we show the kronecker product representation of each record from Table 6. The
573 left column of Table 7 shows the record and its kronecker representation. The next 12 columns show
574 the resulting expansion. Each record becomes as 12-dimensional vector and the column labels in
575 Table 7 show which record is associated with which index in the 12-dimensional vector.

576 The sum of the kron representations of all the records is the data vector \mathbf{x} . It is again a 12-dimensional
577 vector. At each index i , $\mathbf{x}[i]$ is the number of people whose record is associated index i . For example,
578 the 5th component is associated with the record ‘an2’ and there are 2 people with that record. For
579 mathematical convenience, \mathbf{x} is treated as a column vector, but for display purposes, in Table 7 it is
written as a row vector.

	$ay1$	$ay2$	$ay3$	$an1$	$an2$	$an3$	$by1$	$by2$	$by3$	$bn1$	$bn2$	$bn3$
an2: $[1, 0] \otimes [0, 1] \otimes [0, 1, 0]$	0	0	0	0	1	0	0	0	0	0	0	0
bn3: $[0, 1] \otimes [0, 1] \otimes [0, 0, 1]$	0	0	0	0	0	0	0	0	0	0	0	1
by3: $[0, 1] \otimes [1, 0] \otimes [0, 0, 3]$	0	0	0	0	0	0	0	0	1	0	0	0
an2: $[1, 0] \otimes [0, 1] \otimes [0, 1, 0]$	0	0	0	0	1	0	0	0	0	0	0	0
by3: $[0, 1] \otimes [1, 0] \otimes [0, 0, 3]$	0	0	0	0	0	0	0	0	1	0	0	0
Vector of counts \mathbf{x} :	0	0	0	0	2	0	0	0	2	0	0	1

Table 7: Kron product representations of each record and the whole dataset \mathbf{x} . Nonzero components are shown in bold red.

580

581 B.2 The Marginal Workload and its Representation as a Query Matrix.

582 For this example, we set the marginal workload to consist of 3 marginals $Wkload =$
583 $\{\{Att_1\}, \{Att_1, Att_2\}, \{Att_2, Att_3\}\}$.

584 The marginal on attribute set $\mathbf{A} = \{Att_1\}$ has only two cells, which correspond to the number of
585 people with $Att_1 = a$ (i.e., 3) and the number with $Att_1 = b$ (i.e., 3). This is called a one-way

$\mathbf{A} = \{Att_1\}$	$\mathbf{A} = \{Att_1, Att_2\}$	$\mathbf{A} = \{Att_2, Att_3\}$												
	$\mathbf{y} \quad \mathbf{n}$	$\mathbf{1} \quad \mathbf{2} \quad \mathbf{3}$												
\mathbf{a} <table border="1" style="display: inline-table; vertical-align: middle;"><tr><td style="padding: 2px 10px;">2</td></tr><tr><td style="padding: 2px 10px;">3</td></tr></table>	2	3	\mathbf{a} <table border="1" style="display: inline-table; vertical-align: middle;"><tr><td style="padding: 2px 10px;">0</td><td style="padding: 2px 10px;">2</td></tr><tr><td style="padding: 2px 10px;">2</td><td style="padding: 2px 10px;">1</td></tr></table>	0	2	2	1	\mathbf{y} <table border="1" style="display: inline-table; vertical-align: middle;"><tr><td style="padding: 2px 10px;">0</td><td style="padding: 2px 10px;">0</td><td style="padding: 2px 10px;">2</td></tr><tr><td style="padding: 2px 10px;">0</td><td style="padding: 2px 10px;">2</td><td style="padding: 2px 10px;">1</td></tr></table>	0	0	2	0	2	1
2														
3														
0	2													
2	1													
0	0	2												
0	2	1												
\mathbf{b}	\mathbf{b}	\mathbf{n}												

Table 8: True answers to the marginal queries in the marginal workload $Wkload = \{\{Att_1\}, \{Att_1, Att_2\}, \{Att_2, Att_3\}\}$.

586 marginal. The other marginals are two-way marginals because they involve two attributes. For
 587 example, the marginal on $\mathbf{A} = \{Att_2, Att_3\}$ has 6 cells. It represents the number of people for
 588 each combination of values for Att_2 and Att_3 . For example, there are 2 people with $Att_2 = y$ and
 589 $Att_3 = 3$.

590 For each set \mathbf{A} , the marginal on those attributes can be represented as a matrix $\mathbf{Q}_{\mathbf{A}}$ such that
 591 calculating the marginal is equivalent to the matrix-vector multiplication $\mathbf{Q}_{\mathbf{A}}\mathbf{x}$. The construction of
 592 the matrix $\mathbf{Q}_{\mathbf{A}}$ is straightforward. It is a kronecker product of 3 matrices. Each matrix corresponds to
 593 an attribute. If the attribute is in \mathbf{A} then the corresponding term is the identity matrix, otherwise is is
 594 the row vector full of ones. For example, $\mathbf{Q}_{\{Att_1\}}$ is a kron product of 3 matrices: the first matrix
 595 corresponds to Att_1 and is the 2×2 identity matrix. The second matrix is actually the vector full
 596 of ones because Att_2 is not part of the marginal. This vector has 2 components because Att_2 has 2
 597 possible values. Similarly, the third matrix is the vector full of ones with 3 components.

598 For the marginals in $Wkload$, these are the the corresponding matrices:

$$\begin{aligned}
 \mathbf{Q}_{\{Att_1\}} &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \otimes [1 \ 1] \otimes [1 \ 1 \ 1] \\
 &= \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix} \\
 \mathbf{Q}_{\{Att_1, Att_2\}} &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \otimes \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \otimes [1 \ 1 \ 1] \\
 &= \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \end{bmatrix} \\
 \mathbf{Q}_{\{Att_2, Att_3\}} &= [1 \ 1] \otimes \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \otimes \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \\
 &= \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}
 \end{aligned}$$

599 If we multiply $\mathbf{Q}_{\{attr_2, Att_3\}}$ by the data vector \mathbf{x} from Table 7, we get:

$$\mathbf{Q}_{\{attr_2, Att_3\}}\mathbf{x} = \begin{bmatrix} 0 \\ 0 \\ 2 \\ 0 \\ 2 \\ 1 \end{bmatrix}$$

600 Comparing it to the marginals shown in Table 8 we see that it is the flattened version of the marginal.
 601 That is, we take the first column of the $\{Att_2, Att_3\}$ marginal of Table 8, then we put the next column
 602 below it, and the third column is placed at the bottom.

603 **B.3 The Base Mechanisms**

Recall that our workload, the marginals we want privacy-preserving answers to, is $Wkload = \{\{Att_1\}, \{Att_1, Att_2\}, \{Att_2, Att_3\}\}$. Its closure, denoted by $\text{closure}(Wkload)$ is all of its subsets. So,

$$\text{closure}(Wkload) = \{\emptyset, \{Att_1\}, \{Att_2\}, \{Att_3\}, \{Att_1, Att_2\}, \{Att_2, Att_3\}\}$$

604 For each $\mathbf{A} \in \text{closure}(Wkload)$ we need to form a base mechanism $\mathcal{M}_{\mathbf{A}}$. Each $\mathcal{M}_{\mathbf{A}}$ has a free
 605 parameter $\sigma_{\mathbf{A}}^2$ that we are free to choose. Each mechanism $\mathcal{M}_{\mathbf{A}}$ has the form $\mathcal{M}_{\mathbf{A}}(\mathbf{x}; \sigma_{\mathbf{A}}^2) =$
 606 $\mathbf{R}_{\mathbf{A}}\mathbf{x} + N(\mathbf{0}, \Sigma_{\mathbf{A}})$. That is, on input \mathbf{x} , the mechanism multiplies it by a special “residual” matrix
 607 $\mathbf{R}_{\mathbf{A}}$ and then adds correlated Gaussian noise, with zero mean and with covariance matrix $\sigma_{\mathbf{A}}^2 \Sigma_{\mathbf{A}}$.
 608 The residual and covariance matrices for each base mechanism are shown below.

$$\begin{aligned} \mathcal{M}_{\emptyset} : \quad \mathbf{R}_{\emptyset} &= [1 \ 1] \otimes [1 \ 1] \otimes [1 \ 1 \ 1] \\ \Sigma_{\emptyset} &= [1] \end{aligned}$$

$$\begin{aligned} \mathcal{M}_{\{Att_1\}} : \quad \mathbf{R}_{\{Att_1\}} &= [1 \ -1] \otimes [1 \ 1] \otimes [1 \ 1 \ 1] \\ \Sigma_{\{Att_1\}} &= [1 \ -1] ([1 \ -1])^T = [2] \end{aligned}$$

$$\begin{aligned} \mathcal{M}_{\{Att_2\}} : \quad \mathbf{R}_{\{Att_2\}} &= [1 \ 1] \otimes [1 \ -1] \otimes [1 \ 1 \ 1] \\ \Sigma_{\{Att_2\}} &= [1 \ -1] ([1 \ -1])^T = [2] \end{aligned}$$

$$\begin{aligned} \mathcal{M}_{\{Att_3\}} : \quad \mathbf{R}_{\{Att_3\}} &= [1 \ 1] \otimes [1 \ 1] \otimes \begin{bmatrix} 1 & -1 & 0 \\ 1 & 0 & -1 \end{bmatrix} \\ \Sigma_{\{Att_3\}} &= \begin{bmatrix} 1 & -1 & 0 \\ 1 & 0 & -1 \end{bmatrix} \left(\begin{bmatrix} 1 & -1 & 0 \\ 1 & 0 & -1 \end{bmatrix} \right)^T = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix} \end{aligned}$$

$$\begin{aligned} \mathcal{M}_{\{Att_1, Att_2\}} : \quad \mathbf{R}_{\{Att_1, Att_2\}} &= [1 \ -1] \otimes [1 \ -1] \otimes [1 \ 1 \ 1] \\ \Sigma_{\{Att_1, Att_2\}} &= \left([1 \ -1] \otimes [1 \ -1] \right) \left([1 \ -1] \otimes [1 \ -1] \right)^T = [4] \end{aligned}$$

$$\begin{aligned} \mathcal{M}_{\{Att_2, Att_3\}} : \quad \mathbf{R}_{\{Att_2, Att_3\}} &= [1 \ 1] \otimes [1 \ -1] \otimes \begin{bmatrix} 1 & -1 & 0 \\ 1 & 0 & -1 \end{bmatrix} \\ \Sigma_{\{Att_2, Att_3\}} &= \left([1 \ -1] \otimes \begin{bmatrix} 1 & -1 & 0 \\ 1 & 0 & -1 \end{bmatrix} \right) \left([1 \ -1] \otimes \begin{bmatrix} 1 & -1 & 0 \\ 1 & 0 & -1 \end{bmatrix} \right)^T \\ &= \begin{bmatrix} 4 & 2 \\ 2 & 4 \end{bmatrix} \end{aligned}$$

609 Note that for any \mathbf{A} , the residual matrix $\mathbf{R}_{\mathbf{A}}$ has a similar structure to $\mathbf{Q}_{\mathbf{A}}$ except that where $\mathbf{Q}_{\mathbf{A}}$
 610 has an identity matrix in its kron product, $\mathbf{R}_{\mathbf{A}}$ has a subtraction matrix (e.g. $[1 \ -1]$ or $\begin{bmatrix} 1 & -1 & 0 \\ 1 & 0 & -1 \end{bmatrix}$).
 611 Meanwhile the covariance matrix $\Sigma_{\mathbf{A}}$ looks like $\mathbf{R}_{\mathbf{A}}\mathbf{R}_{\mathbf{A}}^T$ except that the vectors full of 1s have been
 612 first removed.

613 How do we interpret the residual matrices? Well, \mathbf{R}_{\emptyset} is the sum query. In fact the matrix vector
 614 multiplication $\mathbf{R}_{\emptyset}\mathbf{x}$ gives us the total number of people in the data.

615 Next, $\mathbf{R}_{\{Att_1\}}$ tells us the information contained in the marginal on $\{Att_1\}$ that is not contained in
 616 the sum query. If we know the total number of people in the data, then the only new information the
 617 marginal gives us is the difference between the number of people with $Att_1 = a$ and the number of
 618 people with $Att_1 = b$. In other words, $\mathbf{R}_{\{Att_1\}}\mathbf{x}$ is this difference. Given this difference, and the
 619 total, once can recover the marginal on attribute Att_1 .

620 Similarly, $\mathbf{R}_{\{Att_2\}}$ contains the information in the marginal on $\{Att_2\}$ that is not provided by the
 621 sum query. Finally $\mathbf{R}_{\{Att_3\}}$ contains the information in the marginal on $\{Att_3\}$ not provided in the

622 sum query, which is the number of people with $Att_3 = 1$ minus the number with $Att_3 = 2$, and
 623 also the number of people with $Att_3 = 1$ minus the number with $Att_3 = 3$. The product $\mathbf{R}_{\{Att_3\}\mathbf{x}}$
 624 returns those two differences as a vector with two components.

625 Now, $\mathbf{R}_{\{Att_1, Att_2\}}$ and $\mathbf{R}_{\{Att_2, Att_3\}}$ are more complicated, but have the same idea. For example,
 626 $\mathbf{R}_{\{Att_1, Att_2\}}$ represents new information that the marginal on $\{Att_1, Att_2\}$ provides that is not
 627 captured by the sub-marginals (the marginal on $\{Att_1\}$ and the marginal on $\{Att_2\}$).

628 In general, the matrix $\mathbf{R}_{\mathbf{A}}$ represents the new information on that the marginal on \mathbf{A} provides, which
 629 is not captured by the marginals on \mathbf{A}' , for $\mathbf{A}' \subset \mathbf{A}$ (strict subsets).

630 Now, Theorem 4.2 tells us that if we take all of the rows of all of the residual matrices, they will be
 631 linearly independent. Furthermore, given an attribute set \mathbf{A} , the total number of rows of $\mathbf{R}_{\mathbf{A}'}$ for all
 632 $\mathbf{A}' \subseteq \mathbf{A}$ is the number of rows in $\mathbf{Q}_{\mathbf{A}}$. Furthermore, the space spanned by those rows is the same as
 633 the space spanned by the rows of $\mathbf{Q}_{\mathbf{A}}$.

634 This also means that if we know $\mathbf{R}_{\mathbf{A}'\mathbf{x}}$ for all $\mathbf{A}' \subseteq \mathbf{A}$ then we can figure out $\mathbf{Q}_{\mathbf{A}}\mathbf{x}$ (and vice versa).

635 Now, we want to get privacy-preserving (noisy) answers to the marginal queries in $Wkload =$
 636 $\{\{Att_1\}, \{Att_1, Att_2\}, \{Att_2, Att_3\}\}$ that are as accurate as possible subject to privacy constraints.
 637 We quantify accuracy using a regular (Definition 4.1) loss function (e.g., sum of the variances of
 638 the answers to the marginals) and we quantify privacy by setting privacy parameters for either
 639 (ϵ, δ) -differential privacy, ρ -zCDP, or μ -Gaussian differential privacy.

640 Theorem 4.4 says that to maximize the accuracy subject to privacy con-
 641 straints, we need to take the closure of the workload, $\text{closure}(Wkload) =$
 642 $\{\emptyset, \{Att_1\}, \{Att_2\}, \{Att_3\}, \{Att_1, Att_2\}, \{Att_2, Att_3\}\}$ and carefully choose positive
 643 numbers $\sigma_{\mathbf{A}}^2$ for each $\mathbf{A} \in \text{closure}(Wkload)$ – so that is 6 numbers total. These numbers are chosen
 644 without looking at the data (we explain how in Section B.6). Once we have these numbers, we run
 645 the mechanisms $\mathcal{M}_{\mathbf{A}}(\mathbf{x}; \sigma_{\mathbf{A}}^2)$ and return their outputs. In other words, we must release the outputs of:

- 646 • $\mathcal{M}_{\emptyset}(\mathbf{x}; \sigma_{\emptyset}^2)$ – produces 1 number (a vector with just one component)
- 647 • $\mathcal{M}_{\{Att_1\}}(\mathbf{x}; \sigma_{\{Att_1\}}^2)$ – produces 1 number (a vector with just one component)
- 648 • $\mathcal{M}_{\{Att_2\}}(\mathbf{x}; \sigma_{\{Att_2\}}^2)$ – produces 1 number (a vector with just one component)
- 649 • $\mathcal{M}_{\{Att_3\}}(\mathbf{x}; \sigma_{\{Att_3\}}^2)$ – produces 2 numbers (a vector with 2 components)
- 650 • $\mathcal{M}_{\{Att_1, Att_2\}}(\mathbf{x}; \sigma_{\{Att_1, Att_2\}}^2)$ – produces 1 number (a vector with 1 component)
- 651 • $\mathcal{M}_{\{Att_2, Att_3\}}(\mathbf{x}; \sigma_{\{Att_2, Att_3\}}^2)$ – produces 2 numbers (a vector with 2 components)

652 Which gives us 8 total (noisy) numbers. In fact, any matrix mechanism for this workload must return
 653 at least 8 noisy numbers, by Theorem 4.4.

654 From these outputs, one can reconstruct noisy answers to the marginals in $Wkload$ (actually one can
 655 reconstruct noisy answers to any marginal in $\text{closure}(Wkload)$). We show how to do this in Section
 656 B.4. Then we show how to compute the privacy cost and variances of the algorithm in Section B.5.

657 B.4 Reconstruction

658 Let $\omega_{\mathbf{A}}$ denote the output of $\mathcal{M}_{\mathbf{A}}$. Thus, after running

- 659 • $\mathcal{M}_{\emptyset}(\mathbf{x}; \sigma_{\emptyset}^2)$
- 660 • $\mathcal{M}_{\{Att_1\}}(\mathbf{x}; \sigma_{\{Att_1\}}^2)$
- 661 • $\mathcal{M}_{\{Att_2\}}(\mathbf{x}; \sigma_{\{Att_2\}}^2)$
- 662 • $\mathcal{M}_{\{Att_3\}}(\mathbf{x}; \sigma_{\{Att_3\}}^2)$
- 663 • $\mathcal{M}_{\{Att_1, Att_2\}}(\mathbf{x}; \sigma_{\{Att_1, Att_2\}}^2)$
- 664 • and $\mathcal{M}_{\{Att_2, Att_3\}}(\mathbf{x}; \sigma_{\{Att_2, Att_3\}}^2)$

665 we have the noisy answers

$$666 \omega_\emptyset, \omega_{\{Att_1\}}, \omega_{\{Att_2\}}, \omega_{\{Att_3\}}, \omega_{\{Att_1, Att_2\}}, \omega_{\{Att_2, Att_3\}}$$

667 From these noisy answers we can produce noisy answers for any marginal in $Wkload$ or even
 668 closure($Wkload$). To reconstruct a marginal on \mathbf{A} , we need $\omega_{\mathbf{A}'}$ for all $\mathbf{A}' \subseteq \mathbf{A}$ – this is not a lot as
 669 these vectors represent as many noisy numbers as there are cells in the desired histogram. So, for
 670 example, if we want to get noisy answers for the marginal on $\{Att_2, Att_3\}$ (which has 6 cells), we
 671 need to use $\omega_\emptyset, \omega_{\{Att_2\}}, \omega_{\{Att_3\}},$ and $\omega_{\{Att_2, Att_3\}}$ (together these ω vectors represent a total of 6
 672 noisy numbers).

In order to reconstruct the marginal on \mathbf{A} , Algorithm 2 multiplies each $\omega_{\mathbf{A}'}$ by a matrix that depends on both \mathbf{A} and \mathbf{A}' . The algorithm calls this matrix \mathbf{U} , but to make the notation precise for this runthrough, we will call it $\mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'}$ (the \mathbf{U} matrix that multiplies $\omega_{\mathbf{A}'}$ when reconstructing \mathbf{A}). It turns out that:

$$\mathbf{Q}_{\mathbf{A}} \mathbf{x} = \sum_{\mathbf{A}' \subseteq \mathbf{A}} \mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'} \mathbf{R}_{\mathbf{A}'} \mathbf{x}$$

673 which means that the marginal on \mathbf{A} could be recreated if we know the quantities $\mathbf{R}_{\mathbf{A}'} \mathbf{x}$ (recall $\mathbf{R}_{\mathbf{A}'}$
 674 are the matrices used to define our base mechanisms). Now, since $\omega_{\mathbf{A}'}$ is a noisy version of $\mathbf{R}_{\mathbf{A}'} \mathbf{x}$,
 675 we can get noisy marginal answers by substituting in these noisy values into the above equation.

676 For example, to reconstruct a noisy answer to the marginal on $\{Att_2, Att_3\}$, we do the following:

$$\begin{aligned} \text{Noisy Marginal on } \{Att_2, Att_3\} &= (\mathbf{U}_{\{Att_2, Att_3\} \leftarrow \emptyset}) \omega_\emptyset \\ &\quad + (\mathbf{U}_{\{Att_2, Att_3\} \leftarrow \{Att_2\}}) \omega_{\{Att_2\}} \\ &\quad + (\mathbf{U}_{\{Att_2, Att_3\} \leftarrow \{Att_3\}}) \omega_{\{Att_3\}} \\ &\quad + (\mathbf{U}_{\{Att_2, Att_3\} \leftarrow \{Att_2, Att_3\}}) \omega_{\{Att_2, Att_3\}} \end{aligned}$$

where

$$\begin{aligned} \mathbf{U}_{\{Att_2, Att_3\} \leftarrow \emptyset} &= \left(\frac{1}{2} \mathbf{1}_2 \right) \otimes \left(\frac{1}{3} \mathbf{1}_3 \right) = \begin{bmatrix} 1/2 \\ 1/2 \end{bmatrix} \otimes \begin{bmatrix} 1/3 \\ 1/3 \\ 1/3 \end{bmatrix} \\ \mathbf{U}_{\{Att_2, Att_3\} \leftarrow \{Att_2\}} &= (\mathbf{Sub}_2^\dagger) \otimes \left(\frac{1}{3} \mathbf{1}_3 \right) = \begin{bmatrix} 1/2 \\ -1/2 \end{bmatrix} \otimes \begin{bmatrix} 1/3 \\ 1/3 \\ 1/3 \end{bmatrix} \\ \mathbf{U}_{\{Att_2, Att_3\} \leftarrow \{Att_3\}} &= \left(\frac{1}{2} \mathbf{1}_2 \right) \otimes (\mathbf{Sub}_3^\dagger) = \begin{bmatrix} 1/2 \\ 1/2 \end{bmatrix} \otimes \begin{bmatrix} 1/3 & 1/3 \\ -2/3 & 1/3 \\ 1/3 & -2/3 \end{bmatrix} \\ \mathbf{U}_{\{Att_2, Att_3\} \leftarrow \{Att_2, Att_3\}} \omega_{\{Att_2, Att_3\}} &= (\mathbf{Sub}_2^\dagger) \otimes (\mathbf{Sub}_3^\dagger) = \begin{bmatrix} 1/2 \\ -1/2 \end{bmatrix} \otimes \begin{bmatrix} 1/3 & 1/3 \\ -2/3 & 1/3 \\ 1/3 & -2/3 \end{bmatrix} \end{aligned}$$

677 Note \mathbf{Sub}_2^\dagger and \mathbf{Sub}_3^\dagger are defined in Lemma 4.6.

678 B.5 Privacy Cost and Marginal Variances

679 Recall that for a marginal workload $Wkload$, we need to run a mechanism $\mathcal{M}_{\mathbf{A}}$ for each $\mathbf{A} \in$
 680 closure($Wkload$). Theorem 4.5 shows how to compute the privacy cost $pcost$ of each. In our
 681 running example, this means:

$$\begin{aligned} 682 \bullet \text{ } pcost(\mathcal{M}_\emptyset(\mathbf{x}; \sigma_\emptyset^2)) &= \frac{1}{\sigma_\emptyset^2} \\ 683 \bullet \text{ } pcost(\mathcal{M}_{\{Att_1\}}(\mathbf{x}; \sigma_{\{Att_1\}}^2)) &= \frac{1}{\sigma_{\{Att_1\}}^2} * \frac{1}{2} \\ 684 \bullet \text{ } pcost(\mathcal{M}_{\{Att_2\}}(\mathbf{x}; \sigma_{\{Att_2\}}^2)) &= \frac{1}{\sigma_{\{Att_2\}}^2} * \frac{1}{2} \\ 685 \bullet \text{ } pcost(\mathcal{M}_{\{Att_3\}}(\mathbf{x}; \sigma_{\{Att_3\}}^2)) &= \frac{1}{\sigma_{\{Att_3\}}^2} * \frac{2}{3} \\ 686 \bullet \text{ } pcost(\mathcal{M}_{\{Att_1, Att_2\}}(\mathbf{x}; \sigma_{\{Att_1, Att_2\}}^2)) &= \frac{1}{\sigma_{\{Att_1, Att_2\}}^2} * \frac{1}{2} * \frac{1}{2} \end{aligned}$$

687 • and $pcost(\mathcal{M}_{\{Att_2, Att_3\}}(\mathbf{x}; \sigma_{\{Att_2, Att_3\}}^2)) = \frac{1}{\sigma_{\{Att_2, Att_3\}}^2} * \frac{1}{2} * \frac{2}{3}$

688 The total privacy cost is,

$$\frac{1}{\sigma_{\emptyset}^2} + \frac{1}{2} \frac{1}{\sigma_{\{Att_1\}}^2} + \frac{1}{2} \frac{1}{\sigma_{\{Att_2\}}^2} + \frac{2}{3} \frac{1}{\sigma_{\{Att_3\}}^2} + \frac{1}{4} \frac{1}{\sigma_{\{Att_1, Att_2\}}^2} + \frac{1}{3} \frac{1}{\sigma_{\{Att_2, Att_3\}}^2}$$

689 Thus this is a symbolic expression in terms of the (currently unknown) noise scale parameters $\sigma_{\mathbf{A}}^2$.
 690 According to Definition 2.3, we can convert the privacy cost to the ρ in ρ -zCDP by dividing by 2 and
 691 we can convert it to the μ from μ -Gaussian DP by taking the square root.

692 For our running example, $Wkload = \{\{Att_1\}, \{Att_1, Att_2\}, \{Att_2, Att_3\}\}$ and we can express the
 693 variance of these marginals (after reconstruction from the noisy $\omega_{\mathbf{A}}$ answers) also in terms of the
 694 noise scale parameters. We do this with the help of Theorem 4.7.

695 • Marginal on $\{Att_1\}$. This marginal is reconstructed from the noisy answers ω_{\emptyset} and $\omega_{\{Att_1\}}$
 696 and so the variance of its cells depends only on σ_{\emptyset}^2 and $\sigma_{\{Att_1\}}^2$. Applying Theorem 4.7, get
 697 that the variance in each cell of this marginal is the same and equals.

$$\left(\sigma_{\emptyset}^2 * \frac{1}{2^2}\right) + \left(\sigma_{\{Att_1\}}^2 * \frac{1}{2}\right)$$

698 • Marginal on $\{Att_1, Att_2\}$. This marginal is reconstructed from $\omega_{\emptyset}, \omega_{\{Att_1\}}, \omega_{\{Att_2\}}$, and
 699 $\omega_{\{Att_1, Att_2\}}$ and hence the variance of the cells in the marginal depend on the corresponding
 700 4 noise scale parameters. The cell variance is

$$\left(\sigma_{\emptyset}^2 * \frac{1}{2^2} * \frac{1}{2^2}\right) + \left(\sigma_{\{Att_1\}}^2 * \frac{1}{2} * \frac{1}{2^2}\right) + \left(\sigma_{\{Att_2\}}^2 * \frac{1}{2} * \frac{1}{2^2}\right) + \left(\sigma_{\{Att_1, Att_2\}}^2 * \frac{1}{2} * \frac{1}{2}\right)$$

701 • Marginal on $\{Att_2, Att_3\}$. Similarly, this marginal also depends on 4 noise scale parameters
 702 as follows:

$$\left(\sigma_{\emptyset}^2 * \frac{1}{2^2} * \frac{1}{3^2}\right) + \left(\sigma_{\{Att_2\}}^2 * \frac{1}{2} * \frac{1}{3^2}\right) + \left(\sigma_{\{Att_3\}}^2 * \frac{2}{3} * \frac{1}{2^2}\right) + \left(\sigma_{\{Att_2, Att_3\}}^2 * \frac{1}{2} * \frac{2}{3}\right)$$

703 B.6 The Sum-of-Variates Loss Function

704 Now we can express the overall privacy cost symbolically in terms of the noise scale parameters.
 705 We can also express the variance of each marginal cell symbolically. We can use these symbolic
 706 expressions to set up any regular loss function and then run it through a convex optimizer to solve it.

707 In this section, we give an example for the weighted sum of variances, which is one of the most
 708 popular loss functions for the matrix mechanism in research settings (mostly because this loss function
 709 is easiest to work with).

710 Each marginal has a weight, which we set to be 1 to avoid introducing more symbols, and the
 711 objective function is computed by adding up the cell variances in a marginal, multiplying by the
 712 weight, and adding up over the workload marginals. The marginal on $\{Att_1\}$ has two cells (so we
 713 multiply the cell variance for this marginal, computed in the previous section, by 2). The marginal on
 714 $\{Att_1, Att_2\}$ has 4 cells, and the marginal on $\{Att_2, Att_3\}$ has 6 cells. Thus, after the dust clears,
 715 the sum of the cell variances across the workload marginals is:

$$= \frac{11}{12} \sigma_{\emptyset}^2 + \frac{3}{2} \sigma_{\{Att_1\}}^2 + \frac{5}{6} \sigma_{\{Att_2\}}^2 + \sigma_{\{Att_3\}}^2 + \sigma_{\{Att_1, Att_2\}}^2 + 2 \sigma_{\{Att_2, Att_3\}}^2$$

716 Thus, we can set up the optimization problem: minimize the sum of variances subject to the privacy
 717 cost (computed in Section B.5) being less than some constant c :

$$\arg \min_{\substack{\sigma_{\emptyset}^2, \sigma_{\{Att_1\}}^2 \\ \sigma_{\{Att_2\}}^2, \sigma_{\{Att_3\}}^2 \\ \sigma_{\{Att_1, Att_2\}}^2, \sigma_{\{Att_2, Att_3\}}^2}} \frac{11}{12} \sigma_{\emptyset}^2 + \frac{3}{2} \sigma_{\{Att_1\}}^2 + \frac{5}{6} \sigma_{\{Att_2\}}^2 + \sigma_{\{Att_3\}}^2 + \sigma_{\{Att_1, Att_2\}}^2 + 2 \sigma_{\{Att_2, Att_3\}}^2$$

$$\text{such that } \frac{1}{\sigma_\emptyset^2} + \frac{1}{2} \frac{1}{\sigma_{\{Att_1\}}^2} + \frac{1}{2} \frac{1}{\sigma_{\{Att_2\}}^2} + \frac{2}{3} \frac{1}{\sigma_{\{Att_3\}}^2} + \frac{1}{4} \frac{1}{\sigma_{\{Att_1, Att_2\}}^2} + \frac{1}{3} \frac{1}{\sigma_{\{Att_2, Att_3\}}^2} \leq c$$

718 If we let the coefficient of $\sigma_{\mathbf{A}}$ be denoted by $v_{\mathbf{A}}$ and the coefficient of $1/\sigma_{\mathbf{A}}^2$ be denoted by $p_{\mathbf{A}}$, then
 719 this optimization problem can be written as:

$$\begin{aligned} \sigma_{\mathbf{A}}^2 : \mathbf{A} \in \text{closure}(Wkload) \quad & \arg \min \sum_{\mathbf{A} \in \text{closure}(Wkload)} v_{\mathbf{A}} \sigma_{\mathbf{A}}^2 \\ \text{s.t.} \quad & \sum_{\mathbf{A} \in \text{closure}(Wkload)} \frac{p_{\mathbf{A}}}{\sigma_{\mathbf{A}}^2} \leq c \end{aligned}$$

720 Lemma G.1 in Section G shows that the optimal solution is obtained by computing:

$$\begin{aligned} T &= \left(\sum_{\mathbf{A}} \sqrt{v_{\mathbf{A}} p_{\mathbf{A}}} \right)^2 / c = \left(\sqrt{\frac{11}{12}} * 1 + \sqrt{\frac{3}{2}} * \frac{1}{2} + \sqrt{\frac{5}{6}} * \frac{1}{2} + \sqrt{2/3} + \sqrt{1/4} + \sqrt{2/3} \right)^2 / c \\ &\approx 21.18/c \\ \sigma_{\mathbf{A}}^2 &= \sqrt{T p_{\mathbf{A}} / (c v_{\mathbf{A}})} \approx \sqrt{21.18 p_{\mathbf{A}} / v_{\mathbf{A}} / c} \\ \sigma_\emptyset^2 &\approx \sqrt{21.18 * 12 / 11 / c} \approx 4.8/c \end{aligned}$$

721 etc.

722 C Optimality Proof of ResidualPlanner

723 In this section, we prove the optimality of ResidualPlanner. It takes advantage of the symmetry
 724 inherent in marginals and regular loss functions.

725 The proof sketch is the following. Given one optimal mechanism \mathcal{M} , we can create a variation $\widetilde{\mathcal{M}}$ of
 726 that does the following. (1) $\widetilde{\mathcal{M}}$ modifies each input record by applying some invertible function f_i to
 727 each attribute Att_i (for example, if Att_i is a tertiary attribute, we can modify the value of Att_i for
 728 each record using a function f_i where $f_i(1) = 3, f_i(2) = 1, f_i(3) = 2$). This step can be viewed as
 729 simply renaming the attribute values within an attribute. (2) Then $\widetilde{\mathcal{M}}$ runs \mathcal{M} on the resulting dataset.
 730 Note that marginals can be reconstructed from the output of $\widetilde{\mathcal{M}}$ by first running the reconstruction
 731 one would do for \mathcal{M} and then inverting the f_i functions on the resulting marginals (i.e., rearranging
 732 the cells in each marginal to undo the within-attribute renaming caused by the f_i). This variation $\widetilde{\mathcal{M}}$
 733 has the same privacy properties as \mathcal{M} and the same loss (due to the regularity condition on the loss).
 734 Hence $\widetilde{\mathcal{M}}$ is also optimal. Then we create yet another optimal privacy mechanism \mathcal{M}^* that splits
 735 the privacy budget across all variations of \mathcal{M} and returns their outputs. It turns out that the privacy
 736 cost matrix of \mathcal{M}^* has eigenvectors that are equal to the rows of the residual matrices $\mathbf{R}_{\mathbf{A}}$ used by
 737 ResidualPlanner. Rewriting the privacy cost matrix of \mathcal{M}^* using this eigendecomposition, we create
 738 another mechanism (the mechanism that runs the base mechanisms of ResidualPlanner) that has the
 739 same privacy cost matrix and the same value for the loss and hence is optimal.

740 The rest of this section explains these steps in details with formal proofs and running commentary
 741 that helps to better understand the notation and constructs in the proof.

742 C.1 Notation Review

743 We first start with a review of key notation. Recall that a dataset $\mathcal{D} = \{r_1, \dots, r_n\}$ is a collection of
 744 records. Each record r_i contains attributes Att_1, \dots, Att_{n_a} and each attribute Att_j can take values
 745 $a_1^{(j)}, \dots, a_{|Att_j|}^{(j)}$.

746 An attribute value $a_i^{(j)}$ for attribute Att_j can be represented as a vector using one-hot encoding.
 747 Specifically, let $e_i^{(j)}$ be a row vector of size $|Att_j|$ with a one in component i and 0 everywhere else.
 748 In this way, $e_i^{(j)}$ is a representation of $a_i^{(j)}$.

749 A record r with attributes $Att_1 = a_{i_1}^{(1)}, Att_2 = a_{i_2}^{(2)}, \dots, Att_{n_a} = a_{i_{n_a}}^{(n_a)}$ can thus be represented as
 750 the kron product $e_{i_1}^{(1)} \otimes e_{i_2}^{(2)} \otimes \dots \otimes e_{i_{n_a}}^{(n_a)}$. This vector has a 1 in exactly one position and 0s everywhere
 751 else. The position of the 1 is the *index* of record r .

752 Thus, a data vector \mathbf{x} is a vector of integers. The value at index i is the number of times the record
 753 associated with index i appears in \mathcal{D} .

754 C.2 Permutations

755 For each attribute Att_i , let $\Pi^{(i)}$ be the set of permutations on the numbers $1, \dots, |Att_i|$, so that each
 756 $\pi \in \Pi^{(i)}$ can be interpreted as a permutation (or renaming) of the attributes values of Att_i . We can
 757 also view π as a function on vectors of size $|Att_i|$ that permutes their coordinates. That is, the i^{th}
 758 coordinate of a vector \mathbf{y} is the $\pi(i)^{\text{th}}$ coordinate of $\pi(\mathbf{y})$.

759 One can select a permutation for each attribute $\pi^{(1)} \in \Pi^{(1)}, \dots, \pi^{(n_a)} \in \Pi^{(n_a)}$ and use it to de-
 760 fine a permutation over records. This permutation maps a record represented by the kron product
 761 $e_{i_1}^{(1)} \otimes e_{i_2}^{(2)} \otimes \dots \otimes e_{i_{n_a}}^{(n_a)}$ into $\pi^{(1)}(e_{i_1}^{(1)}) \otimes \pi^{(2)}(e_{i_2}^{(2)}) \otimes \dots \otimes \pi^{(n_a)}(e_{i_{n_a}}^{(n_a)})$. We can think of this permu-
 762 tation $\pi = (\pi^{(1)}, \dots, \pi^{(n_a)})$ as a function that independently renames each attribute value in a record.
 763 Thus this permutation can be extended to datavectors \mathbf{x} . The value of \mathbf{x} at the index associated with
 764 record r is the value of $\pi(\mathbf{x})$ at the index associated with record $\pi(r)$. Another way to look at it is
 765 that $\pi(\mathbf{x})$ is the histogram associated with the dataset $\{\pi(r_1), \pi(r_2), \dots, \pi(r_n)\}$. This permutation
 766 can be represented as a permutation matrix \mathbf{W}_π such that $\mathbf{W}_\pi \mathbf{x} = \pi(\mathbf{x})$.

767 We let $\Pi = \Pi^{(1)} \times \dots \times \Pi^{(n_a)}$ be the set of all such permutations. We call this the space of *renaming*
 768 permutations since each $\pi \in \Pi$ renames the values of each attribute separately.

769 Our first result is that permutation does not affect the privacy parameters of a mechanism.

770 LEMMA C.1. *Let $\mathcal{M}(\mathbf{x}) \equiv \mathbf{B}\mathbf{x} + N(\mathbf{0}, \Sigma)$ be a mechanism that satisfies ρ -zCDP, (ϵ, δ) -approximate
 771 DP, and μ -Gaussian DP. Let π be a permutation of the indices of \mathbf{x} and \mathbf{W}_π the corresponding
 772 permutation matrix. Then $\mathcal{M}_\pi(\mathbf{x}) \equiv \mathbf{B}\mathbf{W}_\pi \mathbf{x} + N(\mathbf{0}, \Sigma)$ satisfies ρ -zCDP, (ϵ, δ) -approximate DP,
 773 and μ -Gaussian DP (i.e., with the same privacy parameters).*

774 *Proof.* The privacy cost $pcost(\mathcal{M})$ of \mathcal{M} is the largest diagonal of $\mathbf{B}^T \Sigma^{-1} \mathbf{B}$. The privacy cost
 775 $pcost(\mathcal{M}_\pi)$ of \mathcal{M}_π is the largest diagonal of $\mathbf{W}_\pi^T \mathbf{B}^T \Sigma^{-1} \mathbf{B} \mathbf{W}_\pi$. The effect of \mathbf{W}_π on both sides is
 776 to permute the rows and columns of $\mathbf{B}^T \Sigma^{-1} \mathbf{B}$ in the same way. Thus the diagonals of $\mathbf{B}^T \Sigma^{-1} \mathbf{B}$
 777 and $\mathbf{W}_\pi^T \mathbf{B}^T \Sigma^{-1} \mathbf{B} \mathbf{W}_\pi$ are the same up to permutation and hence \mathcal{M} and \mathcal{M}_π have the same privacy
 778 cost and therefore the same privacy parameters. \square

779 The next result is that a renaming permutation preserves the accuracy of a marginal derived from the
 780 answer to a mechanism.

781 LEMMA C.2. *Let $Wkload = \{\mathbf{A}_1, \dots, \mathbf{A}_k\}$ be a workload on marginals. Let $\mathcal{M}(\mathbf{x}) \equiv \mathbf{B}\mathbf{x} +$
 782 $N(\mathbf{0}, \Sigma)$ be a mechanism whose output can be used to provide unbiased estimates of those marginals.
 783 Let $\pi \in \Pi$ be a renaming permutation and \mathbf{W}_π the corresponding permutation matrix. Define
 784 $\mathcal{M}_\pi(\mathbf{x}) \equiv \mathbf{B}\mathbf{W}_\pi \mathbf{x} + N(\mathbf{0}, \Sigma)$. Then unbiased answers to $Wkload$ can be obtained from the output
 785 of \mathcal{M}_π and for any regular loss function \mathcal{L} (Definition 4.1), $\mathcal{L}(Var(\mathbf{A}_1; \mathcal{M}), \dots, Var(\mathbf{A}_k; \mathcal{M})) =$
 786 $\mathcal{L}(Var(\mathbf{A}_1; \mathcal{M}_\pi), \dots, Var(\mathbf{A}_k; \mathcal{M}_\pi))$*

787 *Proof.* For each set of attributes $\mathbf{A}_i \in Wkload$, let $\mathbf{Q}_{\mathbf{A}_i}$ be the query matrix of the marginal
 788 (i.e., the true marginal is computed as $\mathbf{Q}_{\mathbf{A}_i} \mathbf{x}$). Then the best linear unbiased estimate of
 789 the marginal on \mathbf{A}_i from the output ω of \mathcal{M} is $\mathbf{Q}_{\mathbf{A}_i} (\mathbf{B}^T \Sigma^{-1} \mathbf{B})^\dagger \mathbf{B}^T \Sigma^{-1} \omega$ and $Var(\mathbf{A}_i; \mathcal{M})$
 790 is the diagonal of the covariance matrix of this estimate, which is $\mathbf{Q}_{\mathbf{A}_i} (\mathbf{B}^T \Sigma^{-1} \mathbf{B})^\dagger \mathbf{Q}_{\mathbf{A}_i}^T$.
 791 Meanwhile, the best linear unbiased estimate of the marginal on \mathbf{A}_i from the output ω'
 792 of \mathcal{M}_π is $\mathbf{Q}_{\mathbf{A}_i} (\mathbf{W}_\pi^T \mathbf{B}^T \Sigma^{-1} \mathbf{B} \mathbf{W}_\pi)^\dagger \mathbf{W}_\pi^T \mathbf{B}^T \Sigma^{-1} \omega'$ and $Var(\mathbf{A}_i; \mathcal{M}_\pi)$ is the diagonal of
 793 $\mathbf{Q}_{\mathbf{A}_i} (\mathbf{W}_\pi^T \mathbf{B}^T \Sigma^{-1} \mathbf{B} \mathbf{W}_\pi)^\dagger \mathbf{Q}_{\mathbf{A}_i}^T = \mathbf{Q}_{\mathbf{A}_i} \mathbf{W}_\pi^T (\mathbf{B}^T \Sigma^{-1} \mathbf{B})^\dagger \mathbf{W}_\pi \mathbf{Q}_{\mathbf{A}_i}^T$.

794 We note that $\mathbf{Q}_{\mathbf{A}_i} \mathbf{W}_\pi^T$ is a permutation of the rows of $\mathbf{Q}_{\mathbf{A}_i}$ (computing a marginal on a dataset in
 795 which attribute values within the same attribute are renamed is the same as computing the marginal

796 on the original dataset and then renaming the marginal cells, which is permutation of the output of
 797 the marginal computation).

798 Therefore the diagonals of $\mathbf{Q}_{\mathbf{A}_i}(\mathbf{B}^T \boldsymbol{\Sigma}^{-1} \mathbf{B})^\dagger \mathbf{Q}_{\mathbf{A}_i}^T$ and $\mathbf{Q}_{\mathbf{A}_i}(\mathbf{W}_\pi^T \mathbf{B}^T \boldsymbol{\Sigma}^{-1} \mathbf{B} \mathbf{W}_\pi)^\dagger \mathbf{Q}_{\mathbf{A}_i}^T$ are the same
 799 up to permutation. Hence the vector $\text{Var}(\mathbf{A}_i; \mathcal{M})$ is the same as the vector $\text{Var}(\mathbf{A}_i; \mathcal{M}_\pi)$ up to
 800 permutation of the components, and hence does not affect a regular loss function \mathcal{L} . \square

801 Finally, we show that there exists an optimal mechanism whose privacy cost matrix exhibits symme-
 802 tries defined by the set of permutaitons Π .

803 **LEMMA C.3.** *Let $Wkload = \{\mathbf{A}_1 \dots, \mathbf{A}_k\}$ be a workload of marginal queries. Let \mathcal{L} be a regular
 804 loss function. Let U be the set of all Gaussian linear mechanisms that can provide unbiased answers
 805 to the marginals in the $Wkload$. Let γ be a real number. Then whenever either of the following
 806 optimization problems are feasible,*

$$\begin{aligned} & \min_{\mathcal{M} \in U} pcost(\mathcal{M}) \quad \text{s.t. } \mathcal{L}(\text{Var}(\mathbf{A}_1; \mathcal{M}), \dots, \text{Var}(\mathbf{A}_k; \mathcal{M})) \leq \gamma \\ & \min_{\mathcal{M} \in U} \mathcal{L}(\text{Var}(\mathbf{A}_1; \mathcal{M}), \dots, \text{Var}(\mathbf{A}_k; \mathcal{M})) \quad \text{s.t. } pcost(\mathcal{M}) \leq \gamma \end{aligned}$$

807 *the feasible optimization problem is minimized by some mechanism of the form $\overline{\mathcal{M}}(\mathbf{x}) \equiv \overline{\mathbf{B}}\mathbf{x} +$
 808 $N(\mathbf{0}, \overline{\boldsymbol{\Sigma}})$ whose privacy cost matrix $\boldsymbol{\Gamma} \equiv \overline{\mathbf{B}}^T \overline{\boldsymbol{\Sigma}}^{-1} \overline{\mathbf{B}}$ has the following symmetries: for all renaming
 809 permutations $\pi \in \Pi$ (with \mathbf{W}_π being the associated permutation matrix), we have $\boldsymbol{\Gamma} = \mathbf{W}_\pi^T \boldsymbol{\Gamma} \mathbf{W}_\pi$
 810 (in other words, permuting the rows has no effect as long as the columns are permuted in the same
 811 way).*

812 *Proof.* Let $\mathcal{M}_{opt}(\mathbf{x}) \equiv \mathbf{B}_{opt}\mathbf{x} + N(\mathbf{0}, \boldsymbol{\Sigma}_{opt})$ be an optimal mechanism to one of these problems. It
 813 may not have the required symmetries, but from it we will construct an optimal mechanism that does.

814 For a permutation π (and corresponding permutation matrix \mathbf{W}_π) and a positive number λ , consider
 815 the mechanism $\mathcal{M}_{\pi, \lambda}(\mathbf{x}) \equiv \mathbf{B}_{opt} \mathbf{W}_\pi \mathbf{x} + N(\mathbf{0}, \lambda \boldsymbol{\Sigma}_{opt})$. By Lemma C.2, this mechanism also
 816 answers the marginals in $Wkload$.

817 Now consider the mechanism $\overline{\mathcal{M}}$ which, on input \mathbf{x} outputs the result of $\mathcal{M}_{\pi, |\Pi|}$ for all $\pi \in \Pi$.

818 The query matrix of $\overline{\mathcal{M}}$ is $\overline{\mathbf{B}} = \begin{bmatrix} \mathbf{B}_{opt} \mathbf{W}_{\pi_1} \\ \vdots \\ \mathbf{B}_{opt} \mathbf{W}_{\pi_{|\Pi|}} \end{bmatrix}$ and the covariance matrix $\overline{\boldsymbol{\Sigma}}$ is a block diagonal matrix

819 with the scaled matrix $|\Pi| \boldsymbol{\Sigma}_{opt}$ in each block. Clearly, by Lemma C.2, it also provides unbiased
 820 answers to the marginals in $Wkload$.

821 First, we claim that the $pcost(\overline{\mathcal{M}}) \leq pcost(\mathcal{M}_{opt})$ so that the privacy parameters are at least as good.
 822 Recall $pcost(\overline{\mathcal{M}})$ is the largest diagonal entry of:

$$\overline{\mathbf{B}}^T \overline{\boldsymbol{\Sigma}}^{-1} \overline{\mathbf{B}} = \frac{1}{|\Pi|} \sum_{\pi \in \Pi} \mathbf{W}_\pi^T \mathbf{B}_{opt}^T \boldsymbol{\Sigma}_{opt}^{-1} \mathbf{B}_{opt} \mathbf{W}_\pi, \quad (4)$$

823 Since the privacy cost $pcost(\mathcal{M}_{\pi, 1})$ is the largest diagonal of $\mathbf{W}_\pi^T \mathbf{B}_{opt}^T \boldsymbol{\Sigma}_{opt}^{-1} \mathbf{B}_{opt} \mathbf{W}_\pi$ and equals
 824 $pcost(\mathcal{M}_{opt})$, Equation 4 (and convexity of the max function) shows that the $pcost(\overline{\mathcal{M}}) \leq$
 825 $pcost(\mathcal{M}_{opt})$.

826 Next we consider the loss function. Let $\mathbf{A}_i \in Wkload$ be a set of attributes and let $\mathbf{Q}_{\mathbf{A}_i}$ be the
 827 corresponding query matrix for the marginal on \mathbf{A}_i . Then the reconstructed variances of the answers
 828 to this marginal, based on the output of $\overline{\mathcal{M}}$ is:

$$\begin{aligned} \text{Var}(\mathbf{A}_i; \overline{\mathcal{M}}) &= \text{diag} \left(\mathbf{Q}_{\mathbf{A}_i} (\overline{\mathbf{B}}^T \overline{\boldsymbol{\Sigma}}^{-1} \overline{\mathbf{B}})^\dagger \mathbf{Q}_{\mathbf{A}_i}^T \right) \\ &= \text{diag} \left(\frac{1}{|\Pi|} \sum_{\pi \in \Pi} \mathbf{Q}_{\mathbf{A}_i} \left(\mathbf{W}_\pi^T \mathbf{B}_{opt}^T \boldsymbol{\Sigma}_{opt}^{-1} \mathbf{B}_{opt} \mathbf{W}_\pi \right)^\dagger \mathbf{Q}_{\mathbf{A}_i}^T \right) \\ &= \frac{1}{|\Pi|} \sum_{\pi \in \Pi} \text{Var}(\mathbf{A}_i; \mathcal{M}_{\pi, 1}) \end{aligned}$$

829 For any $\pi \in \Pi$, Lemma C.2 tells us that $\mathcal{L}(\text{Var}(\mathbf{A}_1; \mathcal{M}_{opt}), \dots, \text{Var}(\mathbf{A}_k; \mathcal{M}_{opt})) =$
830 $\mathcal{L}(\text{Var}(\mathbf{A}_1; \mathcal{M}_{\pi,1}), \dots, \text{Var}(\mathbf{A}_k; \mathcal{M}_{\pi,1}))$ and so regularity of \mathcal{L} (which includes convexity), means
831 that $\mathcal{L}(\text{Var}(\mathbf{A}_1; \overline{\mathcal{M}}), \dots, \text{Var}(\mathbf{A}_k; \overline{\mathcal{M}})) \leq \mathcal{L}(\text{Var}(\mathbf{A}_1; \mathcal{M}_{opt}), \dots, \text{Var}(\mathbf{A}_k; \mathcal{M}_{opt}))$.

832 Thus $\overline{\mathcal{M}}$ is no worse in privacy or utility than \mathcal{M}_{opt} and hence is optimal.

833 Thus we consider the symmetries of the privacy cost matrix of $\overline{\mathcal{M}}$, which is given in Equation 4.
834 Clearly it has the desired symmetry property that $\mathbf{\Gamma} = \mathbf{W}_\pi^T \mathbf{\Gamma} \mathbf{W}_\pi$ for any $\pi \in \Pi$ as the permutation
835 space Π is an algebraic group.

836

□

837 C.3 From permutations to interpretations

838 Let $\mathcal{M}_{opt}(\mathbf{x}) \equiv \mathbf{B}_{opt}\mathbf{x} + N(\mathbf{0}, \mathbf{\Sigma}_{opt})$ be an optimal mechanism that has the symmetries guaranteed
839 by Lemma C.3. Our goal is to use the symmetries in the privacy cost matrix $\mathbf{\Gamma}_{opt} \equiv \mathbf{B}_{opt}^T \mathbf{\Sigma}_{opt}^{-1} \mathbf{B}_{opt}$
840 to examine the structure of $\mathbf{\Gamma}_{opt}$.

841 If $\gamma_{i,j}$ is the $(i, j)^{\text{th}}$ entry of $\mathbf{\Gamma}_{opt}$ and if there is a renaming permutation that maps r_i (the record
842 associated with index i) to some $r_{i'}$ (at index i') and maps r_j to some $r_{j'}$ then $\gamma_{i,j} = \gamma_{i',j'}$. Note
843 that if r_i and r_j have the same values for attributes Att_1 and Att_2 then $r_{i'}$ and $r_{j'}$ must match on the
844 same attributes because renaming permutations just change the names of values within each attribute.
845 Thus we introduce notation for the set of attributes on which two records match:

846 DEFINITION C.4 (Common Attributes). Define ζ to be the function that takes two records and outputs
847 the set of attributes on which they match. We emphasize that $\zeta(r_i, r_j)$ is a set of attributes, not attribute
848 values.

849 This discussion leads to the following result which characterizes the privacy cost matrix of an optimal
850 mechanism.

851 LEMMA C.5. Under the same conditions as Lemma C.3, there exists an optimal mechanism with a
852 privacy cost matrix $\mathbf{\Gamma}_{opt}$ for which the following holds. In addition to the symmetry guaranteed by
853 Lemma C.3, for every subset of attributes $S \subseteq \{Att_1, \dots, Att_{n_a}\}$, there exists a number c_S such
854 that $\gamma_{i,j}$, the $(i, j)^{\text{th}}$ entry of $\mathbf{\Gamma}_{opt}$, is equal to $c_{\zeta(r_i, r_j)}$. In other words, the $(i, j)^{\text{th}}$ entry is completely
855 determined by the set $\zeta(r_i, r_j)$ (recall r_i the record value associated with index i and r_j is the record
856 value associated with index j).

857 *Proof.* By Lemma C.3, there exists an optimal mechanism with privacy cost matrix $\mathbf{\Gamma}_{opt}$ that is
858 invariant under renaming permutations of its rows as long as the columns are permuted in the same
859 way. Thus if r_i is the record value corresponding to position i and r_j is the record value corresponding
860 to position j , there exists a renaming permutation that maps r_i to some $r_{i'}$ and r_j to some $r_{j'}$ if
861 and only if the attributes on which r_i and r_j match are the same as the attributes on which $r_{i'}$
862 and $r_{j'}$ match each other (in symbols: $\zeta(r_i, r_j) = \zeta(r_{i'}, r_{j'})$). When there exists such a renaming
863 permutation then $\gamma_{i,j} = \gamma_{i',j'}$. Thus the value of $\gamma_{i,j}$ is completely determined by $\zeta(r_i, r_j)$ and the
864 result follows. □

865 From Theorem 4.2, we know that the rows of the matrices of \mathbf{R}_A , for all $\mathbf{A} \subseteq \{Att_1, \dots, Att_{n_a}\}$ are
866 a linearly independent basis for \mathbb{R}^d , where $d = \prod_{i=1}^{n_a} |Att_i|$. Thus we call the rows a *residual basis*.

867 DEFINITION C.6. A row vector \mathbf{v} is a residual basis vector if it is a row in \mathbf{R}_A for some $\mathbf{A} \subseteq$
868 $\{Att_1, \dots, Att_{n_a}\}$.

869 We now provide an interpretation of the residual bases. First, for an attribute Att_ℓ , define the vector
870 $e_{i,j}^{(\ell)}$ to be a vector of length $|Att_\ell|$ such that the element at position i is 1, the element at position j
871 is -1 and everywhere else is 0. In other words, $e_{i,j}^{(\ell)} = e_i^{(\ell)} - e_j^{(\ell)}$ (recall $e_i^{(\ell)}$ is 1 in position i and 0
872 everywhere else and is a one-hot encoding of the attribute $a_i^{(\ell)}$). Now, each element of the residual
873 basis has the form $\mathbf{v}^{(1)} \otimes \dots \otimes \mathbf{v}^{(n_a)}$ where, for each ℓ , $\mathbf{v}^{(\ell)}$ is either the vector $\mathbf{1}_{|Att_\ell|}^T$ or a vector
874 $e_{1,i_\ell}^{(\ell)}$. When the vector for attribute Att_ℓ is the vector $\mathbf{1}_{|Att_\ell|}^T$, we say that all attribute values of Att_ℓ

875 are selected. When the vector for Att_ℓ is $e_{1,i_\ell}^{(\ell)}$, then we say attribute value $a_1^{(\ell)}$ is *positively selected*
876 and $a_{i_\ell}^{(\ell)}$ is *negatively selected* (the other attribute values of Att_ℓ are not selected at all). The attributes
877 for which the kron term is not $\mathbf{1}_{|Att_\ell|}^T$ are called the *discriminative* attributes.

878 As an example of this notation and terminology, consider Table 9. Suppose we have three attributes:
879 Att_1 takes values ‘a’ or ‘b’; Att_2 takes values ‘y’ or ‘n’; Att_3 takes values 1 or 2 or 3.

	ay1	ay2	ay3	an1	an2	an3	by1	by2	by3	bn1	bn2	bn3
bn1: $[0, 1] \otimes [0, 1] \otimes [1, 0, 0]$	0	0	0	0	0	0	0	0	0	1	0	0
$[1, 1] \otimes [1, -1] \otimes [1, -1, 0]$	1	-1	0	-1	1	0	1	-1	0	-1	1	0
$[1, -1] \otimes [1, 1] \otimes [1, 0, -1]$	1	0	-1	1	0	-1	-1	0	1	-1	0	1

Table 9: Kron product representations.

880 In this case, the data vector \mathbf{x} would have 12 components. The first component corresponds to the
881 number of appearances of record “a,y,1” in the dataset, the second component corresponds to record
882 “a,y,2” and so on. The records corresponding to each index of \mathbf{x} are listed in order as the column
883 headings in Table 9. The first row shows the representation of record “b,n,1” which is composed of
884 the second value (b) for Att_1 , the second value (n) for Att_2 and the first value (1) for Att_3 . Hence
885 its kron representation is $[0, 1] \otimes [0, 1] \otimes [1, 0, 0]$ and when the kron product is evaluated, the resulting
886 vector has a 1 in the index corresponding to “bn1” (10th column) and 0 everywhere else.

887 The second and third rows show the expansions of two residual basis vectors $[1, 1] \otimes [1, -1] \otimes [1, -1, 0]$
888 (its discriminative attributes are Att_2 and Att_3) and $[1, -1] \otimes [1, 1] \otimes [1, 0, -1]$ (its discriminative
889 attributes are Att_1 and Att_3). Consider again the kron product $[1, 1] \otimes [1, -1] \otimes [1, -1, 0]$. Note that
890 the first part of the kron product, $[1, 1]$ refers to the first attribute and selects both of its values (sets
891 them to 1). The second part of the kron product $[1, -1]$ refers to the Att_2 and positively selects the
892 first attribute value ‘y’ (sets it to 1) and negatively selected the second attribute value ‘n’ (sets it to
893 -1). The third part is $[1, -1, 0]$ and it positively selects the first attribute value, negatively selects the
894 second, but the third attribute value is not selected at all (i.e., the 3rd position is 0). These attribute
895 selections can help us determine what the kron product looks like when it is expanded as follows. For
896 the residual basis vector $\mathbf{v}^{(1)} \otimes \dots \otimes \mathbf{v}^{(n_a)}$ the value at the index associated with a record r is

- 897 • 0 if r has an attribute whose value is not selected by the residual basis vector’s kron product.
898 In this case we say the residual basis vector assigns a 0 to record r . For example, in the
899 residual basis vector corresponding to kron product $[1, 1] \otimes [1, -1] \otimes [1, -1, 0]$, the third
900 value of the third attribute is not selected. For any record that assigns the attribute value 3 to
901 Att_3 , this residual basis vector assigns a 0 to such a record.
- 902 • 1 if for every attribute, the value assigned to it by r is selected (positively or negatively), and
903 the number of negatively selected attribute values is even. In this case we say the residual
904 basis vector assigns a 1 to record r .
- 905 • -1 if the attribute value for each attribute is selected, and the number of negatively selected
906 attribute values is odd. In this case we say the residual basis vector assigns a -1 to record r .

907 For example, for the residual basis vector $[1, 1] \otimes [1, -1] \otimes [1, -1, 0]$, the attribute value 3 for Att_3
908 is not selected. Hence the value at indices corresponding to records an3,bn3,ay3,by3 are all 0 (see Table
909 9). Next, consider the record an2. The value “a” is positively selected, “n” is negatively selected,
910 and “2” is negatively selected. Hence all attributes are selected and an even number of attributes are
911 negatively selected. Therefore the value at the index associated with an2 is 1. Now for the record by2.
912 The “b” is positively selected, “y” is positively selected, and “2” is negatively selected. Hence there
913 are an odd number of negative selections and so the value at the index associated with by2 is -1.

914 With this discussion and associated notation, we can now show that each residual basis vector is an
915 eigenvector of the optimal privacy cost matrix, and the eigenvalue only depends on which attributes
916 are discriminative.

917 **THEOREM C.7.** *Under the same conditions as Lemma C.3, there exists an optimal mechanism such*
918 *that the eigenvectors of its privacy cost matrix Γ are the residual basis vectors (Definition C.6).*
919 *Furthermore, if two residual basis vectors $\mathbf{v}^{(1)} \otimes \dots \otimes \mathbf{v}^{(n_a)}$ and $\mathbf{w}^{(1)} \otimes \dots \otimes \mathbf{w}^{(n_a)}$ have the same*
920 *discriminative attributes (i.e., for all i , $\mathbf{w}^{(i)} \neq \mathbf{1}_{|Att_i|}^T$ if and only $\mathbf{v}^{(i)} \neq \mathbf{1}_{|Att_i|}^T$) then the two residual*

921 *basis vectors have the same eigenvalues (in other words, all rows of the same residual matrix have*
 922 *the same eigenvalues).*

923 *Proof.* Recall from Definition C.4 that $\zeta(r_i, r_j)$ is the set of attributes on which r_i and r_j are equal.

924 Let Γ be the privacy cost matrix guaranteed by Lemma C.5 with the properties guaranteed by Lemma
 925 C.5, namely that for every subset of attributes $S \subseteq \{Att_1, \dots, Att_{n_a}\}$, there exists a number c_S such
 926 that $\gamma_{i,j}$, the (i, j) th entry of Γ , is equal to $c_{\zeta(r_i, r_j)}$ – the constant associated with the set $\zeta(r_i, r_j)$,
 927 where r_i the record value associated with index i and r_j is the record value associated with index j .

928 Let r_ℓ be a record associated with index ℓ . We consider the dot product between a residual basis vector
 929 $\mathbf{v} = \mathbf{v}^{(1)} \otimes \dots \otimes \mathbf{v}^{(n_a)}$ and the ℓ^{th} row of Γ . Since the entries of the ℓ^{th} row are $c_{\zeta(r_\ell, r_1)}, \dots, c_{\zeta(r_\ell, r_d)}$
 930 and the entries of \mathbf{v} are 0,1,-1, this dot product can be expressed as:

$$\sum_{\substack{r \text{ assigned} \\ \text{value 1 by } \mathbf{v}}} c_{\zeta(r_\ell, r)} - \sum_{\substack{r \text{ assigned} \\ \text{value -1 by } \mathbf{v}}} c_{\zeta(r_\ell, r)} \quad (5)$$

931 We analyze this in three cases.

932 **Case 1: \mathbf{v} assigns a 0 to r_ℓ .** In this case, there is an attribute for which r_ℓ has a value that is not
 933 selected. Without loss of generality, we may assume this is the first attribute Att_1 so that $\mathbf{v}^{(1)} = e_{1,i}$
 934 (the vector with a 1 at the first index and -1 at the i^{th} index for some i and 0 everywhere else) and
 935 the value of Att_1 for r_ℓ is therefore not $a_1^{(1)}$ or $a_i^{(1)}$ (because r_ℓ got assigned 0 by \mathbf{v} due to attribute
 936 Att_1). Now, if a record r appears in the left summation of Equation 5 then its value for Att_1 is either
 937 $a_1^{(1)}$ or $a_i^{(1)}$ and it does not match r_ℓ on the first attribute. But this means that we can transform
 938 r into a record r' by replacing $a_1^{(1)}$ and $a_i^{(1)}$ with each other. This r' would be on the right hand
 939 side of the summation (because we are flipping the sign of the selection by \mathbf{v} of attribute Att_1 in
 940 r'). Furthermore r' also does not match r_ℓ on Att_1 and therefore r matches r_ℓ on exactly the same
 941 attributes as r' matches r_ℓ . Thus $\zeta(r_\ell, r) = \zeta(r_\ell, r')$. Thus the summation term from record r is
 942 cancelled out by r' in Equation 5. Using the same argument, we see that every term in the left
 943 summation is canceled out by a unique term in the right summation, and vice versa. Hence, if \mathbf{v}
 944 assigns a 0 to record r_ℓ (i.e., has a 0 in index ℓ when its kron product representation is expanded)
 945 then the dot product between \mathbf{v} and the ℓ^{th} row of Γ is 0.

946 **Case 2: \mathbf{v} assigns a 1 to r_ℓ .** In this case, every attribute of r_ℓ has a value that is (either positively or
 947 negatively) selected by \mathbf{v} and an even number are negatively selected. Our goal is to show that if
 948 some other record r_t is also assigned a 1 by \mathbf{v} , then the dot product between \mathbf{v} and ℓ^{th} row of Γ is the
 949 same as the dot product between \mathbf{v} and the t^{th} row of Γ . That is, we want to show:

$$\sum_{\substack{r \text{ assigned} \\ \text{value 1 by } \mathbf{v}}} c_{\zeta(r_\ell, r)} - \sum_{\substack{r \text{ assigned} \\ \text{value -1 by } \mathbf{v}}} c_{\zeta(r_\ell, r)} = \sum_{\substack{r \text{ assigned} \\ \text{value 1 by } \mathbf{v}}} c_{\zeta(r_t, r)} - \sum_{\substack{r \text{ assigned} \\ \text{value -1 by } \mathbf{v}}} c_{\zeta(r_t, r)} \quad (6)$$

950 Let S be the set of attributes on which r_ℓ and r_t disagree. Now define a mapping ϕ between records
 951 such that ϕ only modifies attributes in S . For each attribute Att in S , it maps the value that record
 952 r_ℓ has into the value that r_t has and vice versa. (For example, suppose $S = \{Att_1, Att_2\}$ and r_ℓ
 953 has values $a_2^{(1)}$ and $a_3^{(2)}$ for those attributes, respectively, and suppose that r_t has values $a_4^{(1)}$ and
 954 $a_5^{(2)}$ for those attributes. Then ϕ changes $a_2^{(1)}$ in Att_1 to $a_4^{(1)}$ and changes $a_4^{(1)}$ into $a_2^{(1)}$; for Att_2
 955 it changes $a_3^{(2)}$ into $a_5^{(2)}$ and changes $a_5^{(2)}$ into $a_3^{(2)}$. Thus $\phi(r_\ell) = r_t$ and $\phi(r_t) = r_\ell$ and ϕ is its
 956 own inverse. Furthermore, for any record r , $\zeta(r_\ell, r) = \zeta(\phi(r_\ell), \phi(r)) = \zeta(r_t, \phi(r))$ since renaming
 957 attribute values the same way in two records does not affect the set of attributes on which they match
 958 (and the last equality is because $\phi(r_\ell) = r_t$).

959 We next note that since r_t and r_ℓ are both assigned 1 by \mathbf{v} , then they must differ on an even number
 960 of discriminative attributes of \mathbf{v} (if they differ on a discriminative attribute, one must have a value
 961 that is positively selected and the other must have a value that is negatively selected – there cannot be

962 a 0 because r_ℓ and r_t are not assigned a 0 by \mathbf{v}). Therefore, due to its definition, ϕ modifies an even
 963 number of discriminative attributes and therefore for any record r , both r and $\phi(r)$ get assigned the
 964 same value by \mathbf{v} .

965 Putting these facts together, we get:

$$\begin{aligned}
 & \sum_{\substack{r \text{ assigned} \\ \text{value 1 by } \mathbf{v}}} c_{\zeta(r_\ell, r)} - \sum_{\substack{r \text{ assigned} \\ \text{value -1 by } \mathbf{v}}} c_{\zeta(r_\ell, r)} \\
 &= \sum_{\substack{\phi(r) \text{ assigned} \\ \text{value 1 by } \mathbf{v}}} c_{\zeta(r_\ell, r)} - \sum_{\substack{\phi(r) \text{ assigned} \\ \text{value -1 by } \mathbf{v}}} c_{\zeta(r_\ell, r)} \quad \text{since } \phi \text{ doesn't change the summation set} \\
 &= \sum_{\substack{\phi(r) \text{ assigned} \\ \text{value 1 by } \mathbf{v}}} c_{\zeta(\phi(r_\ell), \phi(r))} - \sum_{\substack{\phi(r) \text{ assigned} \\ \text{value -1 by } \mathbf{v}}} c_{\zeta(\phi(r_\ell), \phi(r))} \quad \text{since } \phi \text{ preserves the outcome of } \zeta \\
 &= \sum_{\substack{\phi(r) \text{ assigned} \\ \text{value 1 by } \mathbf{v}}} c_{\zeta(r_t, \phi(r))} - \sum_{\substack{\phi(r) \text{ assigned} \\ \text{value -1 by } \mathbf{v}}} c_{\zeta(r_t, \phi(r))} \quad \text{since } \phi(r_\ell) = r_t \\
 &= \sum_{\substack{r' \text{ assigned} \\ \text{value 1 by } \mathbf{v}}} c_{\zeta(r_t, r')} - \sum_{\substack{r' \text{ assigned} \\ \text{value -1 by } \mathbf{v}}} c_{\zeta(r_t, r')} \quad \text{renaming the summation variable from } \phi(r) \text{ to } r'
 \end{aligned}$$

966 and that proves Equation 6

967 **Case 3: \mathbf{v} assigns a -1 to r_ℓ .** In this case, every attribute of r_ℓ has a value that is (either positively
 968 or negatively) selected by \mathbf{v} and an odd number are negatively selected. Our goal is to show that if
 969 some other record r_t is assigned a 1 by \mathbf{v} , then the dot product between \mathbf{v} and ℓ^{th} row of Γ is the
 970 negative of the dot product between \mathbf{v} and the t^{th} row of Γ . That is, we want to show:

$$\sum_{\substack{r \text{ assigned} \\ \text{value 1 by } \mathbf{v}}} c_{\zeta(r_\ell, r)} - \sum_{\substack{r \text{ assigned} \\ \text{value -1 by } \mathbf{v}}} c_{\zeta(r_\ell, r)} = - \sum_{\substack{r \text{ assigned} \\ \text{value 1 by } \mathbf{v}}} c_{\zeta(r_t, r)} + \sum_{\substack{r \text{ assigned} \\ \text{value -1 by } \mathbf{v}}} c_{\zeta(r_t, r)} \quad (7)$$

971 As in the previous case, we define ϕ in the same way and reasoning as before we see that for any
 972 record r , $\zeta(r_\ell, r) = \zeta(\phi(r_\ell), \phi(r)) = \zeta(r_t, \phi(r))$ and since now ϕ must change an odd number of
 973 discriminative attributes (since r_ℓ and r_t are assigned -1 and 1 by \mathbf{v}) then for any record r , the value
 974 assigned to r by \mathbf{v} is the negative of the value assigned to $\phi(r)$ by \mathbf{v} . Thus we have:

$$\begin{aligned}
 & \sum_{\substack{r \text{ assigned} \\ \text{value 1 by } \mathbf{v}}} c_{\zeta(r_\ell, r)} - \sum_{\substack{r \text{ assigned} \\ \text{value -1 by } \mathbf{v}}} c_{\zeta(r_\ell, r)} \\
 &= \sum_{\substack{\phi(r) \text{ assigned} \\ \text{value -1 by } \mathbf{v}}} c_{\zeta(r_\ell, r)} - \sum_{\substack{\phi(r) \text{ assigned} \\ \text{value +1 by } \mathbf{v}}} c_{\zeta(r_\ell, r)} \quad \text{since } \phi \text{ flips the summation sets} \\
 &= \sum_{\substack{\phi(r) \text{ assigned} \\ \text{value -1 by } \mathbf{v}}} c_{\zeta(\phi(r_\ell), \phi(r))} - \sum_{\substack{\phi(r) \text{ assigned} \\ \text{value +1 by } \mathbf{v}}} c_{\zeta(\phi(r_\ell), \phi(r))} \quad \text{since } \phi \text{ preserves the outcome of } \zeta \\
 &= \sum_{\substack{\phi(r) \text{ assigned} \\ \text{value -1 by } \mathbf{v}}} c_{\zeta(r_t, \phi(r))} - \sum_{\substack{\phi(r) \text{ assigned} \\ \text{value +1 by } \mathbf{v}}} c_{\zeta(r_t, \phi(r))} \quad \text{since } \phi(r_\ell) = r_t \\
 &= \sum_{\substack{r' \text{ assigned} \\ \text{value -1 by } \mathbf{v}}} c_{\zeta(r_t, r')} - \sum_{\substack{r \text{ assigned} \\ \text{value +1 by } \mathbf{v}}} c_{\zeta(r_t, r')} \quad \text{renaming the summation variable from } \phi(r') \text{ to } r'
 \end{aligned}$$

975 and that proves Equation 7.

976 Thus what these 3 cases show us are that there exists some constant β such that:

- 977 • If the i^{th} position of the expansion of \mathbf{v} is 0 (i.e., r_i is assigned 0 by \mathbf{v}), then the i^{th} position
978 of $\Gamma\mathbf{v}$ is also 0 (the dot product between the i^{th} row and \mathbf{v} is 0).
- 979 • If the i^{th} position of the expansion of \mathbf{v} is 1 (i.e., r_i is assigned 1 by \mathbf{v}), then the i^{th} position
980 of $\Gamma\mathbf{v}$ is β (the dot product between the i^{th} row and \mathbf{v} is β).
- 981 • If the i^{th} position of the expansion of \mathbf{v} is -1 (i.e., r_i is assigned -1 by \mathbf{v}), then the i^{th} position
982 of $\Gamma\mathbf{v}$ is $-\beta$ (the dot product between the i^{th} row and \mathbf{v} is $-\beta$).

983 Thus \mathbf{v} is an eigenvector of Γ with eigenvalue β . That proves the first part of the theorem.

984 The next part of the theorem is to show that if two residual basis vectors have the same dis-
985 criminative attributes, then they have the same eigenvalue. So let $\mathbf{v} = \mathbf{v}^{(1)} \otimes \dots \otimes \mathbf{v}^{(n_a)}$ and
986 $\mathbf{w} = \mathbf{w}^{(1)} \otimes \dots \otimes \mathbf{w}^{(n_a)}$ be two residual basis vectors that have the same discriminative attributes.
987 Define a renaming permutation π as follows:

- 988 • For an attribute Att_ℓ that is not discriminative for \mathbf{v} (and hence also not for \mathbf{w}), π does not
989 rename its values (i.e., it acts as the identity for those attribute values).
- 990 • For a discriminative attribute Att_ℓ , let e_{1,i_ℓ} be the kron component for \mathbf{v} (i.e., $\mathbf{v}^{(\ell)} = e_{1,i_\ell}$)
991 and let e_{1,j_ℓ} be the kron component for \mathbf{w} . Note the indices i_ℓ and j_ℓ are not equal to 1. In
992 this case, we make π do the following renamings:

- 993 – $a_{i_\ell} \rightarrow a_{j_\ell}$
- 994 – $a_{j_\ell} \rightarrow a_{i_\ell}$
- 995 – The remaining attribute values are unchanged.

996 By considering which records are assigned 1,-1 and 0 by \mathbf{v} and \mathbf{w} , it is clear that π converts \mathbf{v} into
997 \mathbf{w} (and vice versa). Let \mathbf{W} be the matrix representation of the renaming permutation π , so that
998 $\mathbf{W}\mathbf{v} = \mathbf{w}$ and $\mathbf{W}^T\mathbf{w} = \mathbf{v}$ (a permutation matrix is orthogonal, so its inverse is its transpose). Thus,
999 letting β denote the eigenvalue of \mathbf{v} with respect to Γ , we have:

$$\begin{aligned}
\beta\mathbf{v} &= \Gamma\mathbf{v} \\
&= \Gamma\mathbf{W}^T\mathbf{w} \\
&= \mathbf{W}^T\Gamma\mathbf{W}\mathbf{W}^T\mathbf{w} \quad \text{due to the symmetry from Lemma C.3} \\
&= \mathbf{W}^T\Gamma\mathbf{w},
\end{aligned}$$

1000 since \mathbf{W}^T is the inverse of \mathbf{W} and so

$$\beta\mathbf{w} = \beta\mathbf{W}\mathbf{v} = \mathbf{W}\mathbf{W}^T\Gamma\mathbf{w} = \Gamma\mathbf{w}$$

1001 and thus \mathbf{w} has the same eigenvalue as \mathbf{v} . □

1002 Thus each residual basis matrix $\mathbf{R}_\mathbf{A}$ has a useful property: its rows are linearly independent and are
1003 part of the same eigenspace (linear space of vectors with the same eigenvalue) of the privacy cost
1004 matrix Γ of an optimal mechanism. This allows us to prove the main result:

1005 **THEOREM 4.4.** *Given a marginal workload $Wkload$ and a regular loss function \mathcal{L} , suppose the*
1006 *optimization problem (either Equation 1 or 2) is feasible. Then there exist nonnegative constants*
1007 *$\sigma_\mathbf{A}^2$ for each $\mathbf{A} \in \text{closure}(Wkload)$ (the constants do not depend on the data), such that the optimal*
1008 *linear Gaussian mechanism \mathcal{M}_{opt} releases $\mathcal{M}_\mathbf{A}(\mathbf{x}; \sigma_\mathbf{A}^2)$ for all $\mathbf{A} \in \text{closure}(Wkload)$. Furthermore,*
1009 *any matrix mechanism for this workload must release at least this many noise query answers.*

1010 *Proof of Theorem 4.4.* Let ALL represent $\text{closure}(\{Att_1, \dots, Att_{n_a}\})$ – all possible subsets of at-
 1011 tributes. Theorem C.7 guarantees that there is an optimal mechanism whose privacy cost matrix Γ
 1012 has eigenvectors equal to the rows of the residual matrices. Rows within the same residual matrix
 1013 have the same eigenvalues. Since privacy cost matrices are symmetric positive semidefinite, this
 1014 means that for every $\mathbf{A} \in ALL$, there exists a nonnegative number $\beta_{\mathbf{A}}$ such that:

$$\Gamma \mathbf{R}_{\mathbf{A}}^T = \beta_{\mathbf{A}} \mathbf{R}_{\mathbf{A}}^T$$

1015 By Theorem 3.5 of [47], if two Gaussian linear mechanisms have the same privacy cost matrix then
 1016 each can be obtained by linearly processing the other. Thus they have the same privacy properties
 1017 (under any postprocessing invariant privacy definition) and can be used to answer the same queries
 1018 with the same exact accuracies (under any measure of accuracy). Thus we just need to construct the
 1019 appropriate mechanism having privacy cost matrix Γ .

1020 For each \mathbf{A} , let $\mathbf{Z}_{\mathbf{A}}$ be a matrix with orthonormal rows that span the row space of $\mathbf{R}_{\mathbf{A}}$. Thus the rows
 1021 of $\mathbf{Z}_{\mathbf{A}}$ are also eigenvectors of Γ (having common eigenvalue $\beta_{\mathbf{A}}$) and the rows of $\mathbf{Z}_{\mathbf{A}}$ are orthogonal
 1022 to the rows of $\mathbf{Z}_{\mathbf{A}'}$ for $\mathbf{A} \neq \mathbf{A}'$ (a consequence of Theorem 4.2). Thus the set of rows of the $\mathbf{Z}_{\mathbf{A}}$ for
 1023 all $\mathbf{A} \in ALL$ are a complete list of the eigenvectors of Γ (they are linearly independent and span \mathbb{R}^d).
 1024 Thus the (symmetric positive semidefinite) privacy cost matrix Γ can be expressed as:

$$\Gamma = \sum_{\mathbf{A} \in ALL} \beta_{\mathbf{A}} \mathbf{Z}_{\mathbf{A}}^T \mathbf{Z}_{\mathbf{A}}$$

1025 and one mechanism that achieves this privacy cost matrix is the one that releases $\mathbf{Z}_{\mathbf{A}} \mathbf{x} + N(\mathbf{0}, \frac{1}{\beta_{\mathbf{A}}} \mathcal{I})$
 1026 for each $\mathbf{A} \in ALL$ for which $\beta_{\mathbf{A}} \neq 0$ (i.e., we can drop the eigenvectors with eigenvalue equal to 0
 1027 as they make no difference to the privacy cost matrix).

1028 Now, since the rows of $\mathbf{R}_{\mathbf{A}}$ and $\mathbf{Z}_{\mathbf{A}}$ are independent linear bases of the same subspace, then there
 1029 exists an invertible matrix $\mathbf{Y}_{\mathbf{A}}$ such that $\mathbf{R}_{\mathbf{A}} = \mathbf{Y}_{\mathbf{A}} \mathbf{Z}_{\mathbf{A}}$. Furthermore, $\mathbf{R}_{\mathbf{A}} \mathbf{R}_{\mathbf{A}}^T$ is invertible and
 1030 $\mathbf{Z}_{\mathbf{A}} \mathbf{Z}_{\mathbf{A}}^T = \mathcal{I}$ by orthonormality of its rows. Therefore

$$\begin{aligned} \mathbf{R}_{\mathbf{A}}^T (\mathbf{R}_{\mathbf{A}} \mathbf{R}_{\mathbf{A}}^T)^{-1} \mathbf{R}_{\mathbf{A}} &= \mathbf{Z}_{\mathbf{A}}^T \mathbf{Y}_{\mathbf{A}}^T (\mathbf{Y}_{\mathbf{A}} \mathbf{Z}_{\mathbf{A}} \mathbf{Z}_{\mathbf{A}}^T \mathbf{Y}_{\mathbf{A}}^T)^{-1} \mathbf{Y}_{\mathbf{A}} \mathbf{Z}_{\mathbf{A}} \\ &= \mathbf{Z}_{\mathbf{A}}^T \mathbf{Y}_{\mathbf{A}}^T \mathbf{Y}_{\mathbf{A}}^{-T} (\mathbf{Z}_{\mathbf{A}} \mathbf{Z}_{\mathbf{A}}^T)^{-1} \mathbf{Y}_{\mathbf{A}}^{-1} \mathbf{Y}_{\mathbf{A}} \mathbf{Z}_{\mathbf{A}} \\ &= \mathbf{Z}_{\mathbf{A}}^T (\mathbf{Z}_{\mathbf{A}} \mathbf{Z}_{\mathbf{A}}^T)^{-1} \mathbf{Z}_{\mathbf{A}} \\ &= \mathbf{Z}_{\mathbf{A}}^T \mathbf{Z}_{\mathbf{A}} \quad \text{by orthonormality of the rows of } \mathbf{Z}_{\mathbf{A}} \end{aligned}$$

1031 Thus we have

$$\Gamma = \sum_{\mathbf{A} \in ALL} \beta_{\mathbf{A}} \mathbf{R}_{\mathbf{A}}^T (\mathbf{R}_{\mathbf{A}} \mathbf{R}_{\mathbf{A}}^T)^{-1} \mathbf{R}_{\mathbf{A}}$$

1032 and a mechanism that achieves this privacy cost matrix is the one that releases $\mathbf{R}_{\mathbf{A}} \mathbf{x} +$
 1033 $N(\mathbf{0}, \frac{1}{\beta_{\mathbf{A}}} \mathbf{R}_{\mathbf{A}} \mathbf{R}_{\mathbf{A}}^T)$ for each \mathbf{A} for which $\beta_{\mathbf{A}} \neq 0$.

1034 We next note that each covariance matrices we propose to use, $\Sigma_{\mathbf{A}}$, is proportional to $\mathbf{R}_{\mathbf{A}} \mathbf{R}_{\mathbf{A}}^T$ (they
 1035 are equal up to positive rescaling). If we define the positive constants $\kappa_{\mathbf{A}}$ such that $\mathbf{R}_{\mathbf{A}} \mathbf{R}_{\mathbf{A}}^T = \kappa_{\mathbf{A}} \Sigma_{\mathbf{A}}$
 1036 then we note that the $\sigma_{\mathbf{A}}^2$ in the theorem statement are equal to $\kappa_{\mathbf{A}} / \beta_{\mathbf{A}}$.

1037 Next, we show that the eigenvalues $\beta_{\mathbf{A}} > 0$ for $\mathbf{A} \in \text{closure}(Wkload)$ and 0 otherwise, so that the
 1038 optimal mechanism would not make use of any submechanism $\mathcal{M}_{\mathbf{A}}$ for $\mathbf{A} \notin \text{closure}(Wkload)$.

1039 First, by Theorem 4.2, the rows of $\mathbf{R}_{\mathbf{A}}$, for $\mathbf{A} \in \text{closure}(Wkload)$ form an independent linear
 1040 basis for the space spanned by the rows of the marginals $\mathbf{Q}_{\mathbf{A}}$ for $\mathbf{A} \in Wkload$. If a noisy $\mathbf{R}_{\mathbf{A}} \mathbf{x}$
 1041 is not released for some $\mathbf{A} \in \text{closure}(Wkload)$, then an unbiased noisy answer to at least one
 1042 of the workload marginals could not be computed. Hence, they must all be part of the optimal
 1043 mechanism (and thus, because of linear independence, any mechanism needs to get at least as many
 1044 scalar noisy answers as this). This shows that $\beta_{\mathbf{A}} > 0$ for all $\mathbf{A} \in \text{closure}(Wkload)$. On the
 1045 other hand since the rows of $\mathbf{R}_{\mathbf{A}}$ are orthogonal to the rows of $\mathbf{R}_{\mathbf{A}'}$ for $\mathbf{A} \neq \mathbf{A}'$, getting answers
 1046 to $\mathbf{R}_{\mathbf{A}'} \mathbf{x}$, for $\mathbf{A}' \notin \text{closure}(Wkload)$, cannot help estimate the answers to the marginals $\mathbf{Q}_{\mathbf{A}}$ for
 1047 $\mathbf{A} \in Wkload$ (by Theorem 4.2, $\mathbf{R}_{\mathbf{A}'}$ are orthogonal to the matrices representing these marginals
 1048 when $\mathbf{A}' \notin \text{closure}(Wkload)$). Hence an optimal privacy mechanism cannot waste privacy budget
 1049 on these irrelevant queries. This shows that $\beta_{\mathbf{A}'} = 0$ for $\mathbf{A}' \notin \text{closure}(Wkload)$ and concludes the
 1050 proof. \square

1051 **D The other proofs about base mechanisms**

1052 THEOREM 4.2. Let \mathbf{A} be a set of attributes and let $\mathbf{Q}_{\mathbf{A}}$ be the matrix representation of the marginal
 1053 on \mathbf{A} . Then the rows of the matrices $\mathbf{R}_{\mathbf{A}'}$, for all $\mathbf{A}' \subseteq \mathbf{A}$, form a linearly independent basis of the
 1054 row space of $\mathbf{Q}_{\mathbf{A}}$. Furthermore, if $\mathbf{A}' \neq \mathbf{A}''$ then $\mathbf{R}_{\mathbf{A}'} \mathbf{R}_{\mathbf{A}''}^T = \mathbf{0}$ (they are mutually orthogonal).

1055 *Proof of Theorem 4.2.* Consider two sets $\mathbf{A}' \neq \mathbf{A}''$ and represent there respective residual matrices
 1056 as:

$$\begin{aligned}\mathbf{R}_{\mathbf{A}'} &= \mathbf{V}'_1 \otimes \cdots \otimes \mathbf{V}'_{n_a} \\ \mathbf{R}_{\mathbf{A}''} &= \mathbf{V}''_1 \otimes \cdots \otimes \mathbf{V}''_{n_a} \\ \mathbf{R}_{\mathbf{A}'} \mathbf{R}_{\mathbf{A}''}^T &= (\mathbf{V}'_1 (\mathbf{V}'_1)^T) \otimes \cdots \otimes (\mathbf{V}'_{n_a} (\mathbf{V}'_{n_a})^T)\end{aligned}$$

1057 Since $\mathbf{A}' \neq \mathbf{A}''$ then one of them contains an attribute, say Att_i , that the other doesn't have. Therefore
 1058 either \mathbf{V}'_i or \mathbf{V}''_i is the vector $\mathbf{1}_{|Att_i|}^T$ and the other is $\mathbf{Sub}_{|Att_i|}$. However, $\mathbf{1}_{|Att_i|}^T \mathbf{Sub}_{|Att_i|}^T = \mathbf{0}$ and
 1059 $\mathbf{Sub}_{|Att_i|} \mathbf{1}_{|Att_i|} = \mathbf{0}$ and hence $\mathbf{R}_{\mathbf{A}'} \mathbf{R}_{\mathbf{A}''}^T = \mathbf{0}$.

1060 Next, for any set \mathbf{A}' , it is clear that the row space of $\mathbf{R}_{\mathbf{A}'}$ is contained in the row space of the
 1061 marginal matrix $\mathbf{Q}_{\mathbf{A}'}$. It is also clear that if $\mathbf{A}' \subseteq \mathbf{A}$ then the row space of the marginal matrix $\mathbf{Q}_{\mathbf{A}'}$
 1062 is contained in the row space of $\mathbf{Q}_{\mathbf{A}}$ (because $\mathbf{Q}_{\mathbf{A}'}$ represents a sub-marginal of $\mathbf{Q}_{\mathbf{A}}$). Thus the rows
 1063 of the matrices $\mathbf{R}_{\mathbf{A}'}$, for all $\mathbf{A}' \subseteq \mathbf{A}$, are contained in the rowspace of $\mathbf{Q}_{\mathbf{A}}$. Thus we just need to
 1064 show that the combined rows of $\mathbf{R}_{\mathbf{A}'}$, for all $\mathbf{A}' \subseteq \mathbf{A}$, are linearly independent and that the number
 1065 of rows is the same as the number of rows of $\mathbf{Q}_{\mathbf{A}}$.

1066 First, each $\mathbf{R}_{\mathbf{A}'}$ is a kronecker product of matrices with full row rank, and so $\mathbf{R}_{\mathbf{A}'}$ has full row rank
 1067 (therefore its rows are linearly independent). Furthermore, since $\mathbf{R}_{\mathbf{A}'} \mathbf{R}_{\mathbf{A}''}^T = \mathbf{0}$ whenever $\mathbf{A}' \neq \mathbf{A}''$
 1068 this means that the row space of $\mathbf{R}_{\mathbf{A}'}$ is orthogonal to the row space of $\mathbf{R}_{\mathbf{A}''}$. Hence the combined
 1069 rows of the $\mathbf{R}_{\mathbf{A}'}$, for all $\mathbf{A}' \subseteq \mathbf{A}$, are linearly independent.

1070 Next, the number of rows in \mathbf{R}_{\emptyset} is 1 and the number of rows in $\mathbf{R}_{\mathbf{A}'}$ is equal to $\prod_{Att_i \in \mathbf{A}'} (|Att_i| - 1)$
 1071 for $\mathbf{A}' \neq \emptyset$ and so the total number of rows in the residual matrices is $1 + \sum_{\substack{\mathbf{A}' \subseteq \mathbf{A} \\ \mathbf{A}' \neq \emptyset}} \prod_{Att_i \in \mathbf{A}'} (|Att_i| - 1)$.

1072 By the distributive property of multiplication, this is exactly the same as the product:

$$\prod_{Att_i \in \mathbf{A}} ((|Att_i| - 1) + 1) = \prod_{Att_i \in \mathbf{A}} |Att_i|$$

1073 which is the number of rows in $\mathbf{Q}_{\mathbf{A}}$ and that proves that the combined rows of $\mathbf{R}_{\mathbf{A}'}$, for all $\mathbf{A}' \subseteq \mathbf{A}$,
 1074 form a linearly independent basis for the row span of $\mathbf{Q}_{\mathbf{A}}$. \square

1075 LEMMA D.1. For any i , $\mathbf{Sub}_{|Att_i|}^T (\mathbf{Sub}_{|Att_i|} \mathbf{Sub}_{|Att_i|}^T)^{-1} \mathbf{Sub}_{|Att_i|} = \mathcal{I}_{|Att_i|} - \frac{1}{|Att_i|} \mathbf{1}_{|Att_i|} \mathbf{1}_{|Att_i|}^T$

1076 *Proof of Lemma D.1.* For the moment, let \mathbf{Y} denote $\mathbf{Sub}_{|Att_i|}^T (\mathbf{Sub}_{|Att_i|} \mathbf{Sub}_{|Att_i|}^T)^{-1} \mathbf{Sub}_{|Att_i|}$.
 1077 Then we know:

- 1078 • \mathbf{Y} is symmetric.
- 1079 • \mathbf{Y} is an $|Att_i| \times |Att_i|$ matrix and its rank is $|Att_i| - 1$ since the rank of $\mathbf{Sub}_{|Att_i|}$ is
 1080 $|Att_i| - 1$.
- 1081 • $\mathbf{Sub}_{|Att_i|} \mathbf{Y} \mathbf{Sub}_{|Att_i|}^T = \mathbf{Sub}_{|Att_i|} \mathbf{Sub}_{|Att_i|}^T$.

1082 Now, one symmetric solution to the equation $\mathbf{Sub}_{|Att_i|} \mathbf{X} \mathbf{Sub}_{|Att_i|}^T = \mathbf{Sub}_{|Att_i|} \mathbf{Sub}_{|Att_i|}^T$ is $\mathbf{X} =$
 1083 $\mathcal{I}_{|Att_i|}$ and if \mathbf{X}_1 is another symmetric solution then $\mathbf{Sub}_{|Att_i|} (\mathcal{I}_{|Att_i|} - \mathbf{X}_1) \mathbf{Sub}_{|Att_i|}^T = \mathbf{0}$.

1084 This means that $\mathbf{Sub}_{|Att_i|}\mathbf{v} = \mathbf{0}$ for each eigenvector \mathbf{v} of the symmetric matrix $\mathcal{I}_{|Att_i|} - \mathbf{X}_1$
 1085 that has a nonzero eigenvalue. Since the rank of $\mathbf{Sub}_{|Att_i|}$ is $|Att_i| - 1$, the only vectors \mathbf{v} for
 1086 which $\mathbf{Sub}_{|Att_i|}\mathbf{v} = \mathbf{0}$ are proportional to $\mathbf{1}_{|Att_i|}$ (the null space has rank 1) and so $\mathcal{I}_{|Att_i|} - \mathbf{X}_1 =$
 1087 $-c\mathbf{1}_{|Att_i|}\mathbf{1}_{|Att_i|}^T$ for some constant c .

1088 This means that \mathbf{Y} (and any other symmetric solution) has the form $\mathcal{I}_{|Att_i|} + c\mathbf{1}_{|Att_i|}\mathbf{1}_{|Att_i|}^T$. To find
 1089 c , we note that \mathbf{Y} is not full rank.

1090 By the Sherman-Morrison-Woodbury inversion formula, if $\mathcal{I}_{|Att_i|} + c\mathbf{1}_{|Att_i|}\mathbf{1}_{|Att_i|}^T$ is invertible, then
 1091 its inverse is $\mathcal{I}_{|Att_i|} - c\mathbf{1}_{|Att_i|}\left(1 + c\mathbf{1}_{|Att_i|}^T\mathbf{1}_{|Att_i|}\right)^{-1}\mathbf{1}_{|Att_i|} = \mathcal{I}_{|Att_i|} - c\frac{\mathbf{1}_{|Att_i|}\mathbf{1}_{|Att_i|}^T}{1+c|Att_i|}$. Thus, to
 1092 prevent invertibility, we must have $c = -1/|Att_i|$.

1093 Therefore $\mathbf{Y} = \mathcal{I}_{|Att_i|} - \frac{1}{|Att_i|}\mathbf{1}_{|Att_i|}\mathbf{1}_{|Att_i|}^T$. \square

1094 **THEOREM 4.5.** *The privacy cost of $\mathcal{M}_{\mathbf{A}}$ with noise parameter $\sigma_{\mathbf{A}}^2$ is $\frac{1}{\sigma_{\mathbf{A}}^2} \prod_{Att_i \in \mathbf{A}} \frac{|Att_i| - 1}{|Att_i|}$ and the*
 1095 *evaluation of $\mathcal{M}_{\mathbf{A}}$ given in Algorithm 1 is correct.*

1096 *Proof of Theorem 4.5.* Without loss of generality (and to simplify notation), assume $\mathbf{A} =$
 1097 $\{Att_1, \dots, Att_\ell\}$ consists of the first ℓ attributes.

1098 By definition, $pcost(\mathcal{M}_{\mathbf{A}}(\cdot; \sigma_{\mathbf{A}}^2))$ is the largest diagonal of $\frac{1}{\sigma_{\mathbf{A}}^2} \mathbf{R}_{\mathbf{A}}^T \Sigma_{\mathbf{A}}^{-1} \mathbf{R}_{\mathbf{A}}$. Thus we can write:

$$\begin{aligned}
 \mathbf{R}_{\mathbf{A}} &= \left(\bigotimes_{i=1}^{\ell} \mathbf{Sub}_{|Att_i|} \right) \otimes \left(\bigotimes_{j=\ell+1}^{n_a} \mathbf{1}_{|Att_j|}^T \right) \\
 \mathbf{R}_{\mathbf{A}}^T &= \left(\bigotimes_{i=1}^{\ell} \mathbf{Sub}_{|Att_i|}^T \right) \otimes \left(\bigotimes_{j=\ell+1}^{n_a} \mathbf{1}_{|Att_j|} \right) \\
 \mathbf{H} &= \left(\bigotimes_{i=1}^{\ell} \mathbf{Sub}_{|Att_i|} \right) \otimes \left(\bigotimes_{j=\ell+1}^{n_a} [1] \right) \quad (\text{rightmost krons use } 1 \times 1 \text{ matrices}) \\
 \Sigma_{\mathbf{A}} &= \mathbf{H}\mathbf{H}^T = \left(\bigotimes_{i=1}^{\ell} (\mathbf{Sub}_{|Att_i|} \mathbf{Sub}_{|Att_i|}^T) \right) \otimes \left(\bigotimes_{j=\ell+1}^{n_a} [1] \right) \\
 \Sigma_{\mathbf{A}}^{-1} &= \left(\bigotimes_{i=1}^{\ell} (\mathbf{Sub}_{|Att_i|} \mathbf{Sub}_{|Att_i|}^T)^{-1} \right) \otimes \left(\bigotimes_{j=\ell+1}^{n_a} [1] \right) \\
 \mathbf{R}_{\mathbf{A}}^T \Sigma_{\mathbf{A}}^{-1} \mathbf{R}_{\mathbf{A}} &= \left(\bigotimes_{i=1}^{\ell} \mathbf{Sub}_{|Att_i|}^T (\mathbf{Sub}_{|Att_i|} \mathbf{Sub}_{|Att_i|}^T)^{-1} \mathbf{Sub}_{|Att_i|} \right) \otimes \left(\bigotimes_{j=\ell+1}^{n_a} \mathbf{1}_{|Att_j|} [1] \mathbf{1}_{|Att_j|}^T \right) \tag{8}
 \end{aligned}$$

1099 Now, by Lemma D.1,

$$\mathbf{Sub}_{|Att_i|}^T (\mathbf{Sub}_{|Att_i|} \mathbf{Sub}_{|Att_i|}^T)^{-1} \mathbf{Sub}_{|Att_i|} = \mathcal{I}_{|Att_i|} - \frac{1}{|Att_i|} \mathbf{1}_{|Att_i|} \mathbf{1}_{|Att_i|}^T \tag{9}$$

1100 Since its diagonals are $\frac{|Att_i| - 1}{|Att_i|}$, then combined with Equation 8 it proves the result for
 1101 $pcost(\mathcal{M}_{\mathbf{A}}(\cdot, \sigma_{\mathbf{A}}^2))$.

1102 We next consider the correctness of Algorithm 1. First, we need to show that for the matrix \mathbf{H} defined
 1103 in Line 3 in Algorithm 1, $\mathbf{H}\mathbf{Q}_{\mathbf{A}}\mathbf{x} = \mathbf{R}_{\mathbf{A}}\mathbf{x}$. Then we can write:

$$\mathbf{R}_{\mathbf{A}} = \left(\bigotimes_{i=1}^{\ell} \mathbf{Sub}_{|Att_i|} \right) \otimes \left(\bigotimes_{j=\ell+1}^{n_a} \mathbf{1}_{|Att_j|}^T \right)$$

$$\mathbf{Q}_A = \left(\bigotimes_{i=1}^{\ell} \mathcal{I}_{|Att_i|} \right) \otimes \left(\bigotimes_{j=\ell+1}^{n_a} \mathbf{1}_{|Att_j|}^T \right) \quad \text{rightmost product is a matrix with 1 row}$$

$$\mathbf{H} = \left(\bigotimes_{i=1}^{\ell} \mathbf{Sub}_{|Att_i|} \right) \otimes [1] \quad (\text{rightmost term is a } 1 \times 1 \text{ matrix})$$

$$\begin{aligned} \mathbf{HQ}_A &= \left(\bigotimes_{i=1}^{\ell} (\mathbf{Sub}_{|Att_i|} \mathcal{I}_{|Att_i|}) \right) \otimes \left([1] \left(\bigotimes_{j=\ell+1}^{n_a} \mathbf{1}_{|Att_j|}^T \right) \right) \\ &= \mathbf{R}_A \end{aligned}$$

1104 Next, we note that if \mathbf{z} is distributed as $N(0, \mathbf{I}_m)$ (Line 4 in Algorithm 1) then $\sigma_A \mathbf{H}\mathbf{z}$ has the
1105 distribution $N(0, \sigma^2 \mathbf{H}\mathbf{H}^T) = \Sigma_A$ and hence the algorithm is correct. \square

1106 E Proofs related to the reconstruction step

1107 LEMMA 4.6. For any Att_i , let $\ell = |Att_i|$. The matrix \mathbf{Sub}_ℓ has the following block matrix, with
1108 dimensions $\ell \times (\ell - 1)$, as its pseudo-inverse (and right inverse): $\mathbf{Sub}_\ell^\dagger = \frac{1}{\ell} \begin{bmatrix} \mathbf{1}_{\ell-1}^T \\ \mathbf{1}_{\ell-1} \mathbf{1}_{\ell-1}^T - \ell \mathcal{I}_{\ell-1} \end{bmatrix}$.

1109 *Proof of Lemma 4.6.* First, if a matrix has a right inverse then that is the pseudo-inverse. Hence we
1110 just need to show that $\mathbf{Sub}_\ell \mathbf{Sub}_\ell^\dagger = \mathcal{I}_{\ell-1}$.

1111 Note that the j^{th} row of \mathbf{Sub}_ℓ has a 1 in position 1, -1 in position $j + 1$, and is 0 everywhere else.

1112 Meanwhile, the i^{th} column of our claimed representation of $\mathbf{Sub}_\ell^\dagger$ has a $-(\ell - 1)/\ell$ in position $i + 1$
1113 and $1/\ell$ everywhere else.

1114 Hence if $j \neq i$ then the dot product between row j of \mathbf{Sub}_ℓ and column i of $\mathbf{Sub}_\ell^\dagger$ is 0 since the
1115 nonzero elements of the row from \mathbf{Sub}_ℓ are being multiplied by $1/\ell$ and $1/\ell$.

1116 If $i = j$ then the corresponding first elements that are multiplied are 1 and $1/\ell$ while the elements at
1117 position $i + 1$ being multiplied are -1 and $-(\ell - 1)/\ell$. Furthermore, $1(1/\ell) + (-1)(-(\ell - 1)/\ell) =$
1118 1. \square

1119 LEMMA E.1. For any attribute Att_i , let $\ell = |Att_i|$. Then $\mathbf{Sub}_\ell^\dagger (\mathbf{Sub}_\ell \mathbf{Sub}_\ell^T) \mathbf{Sub}_\ell^{\dagger T} = \mathcal{I}_\ell - \frac{1}{\ell} \mathbf{1}_\ell \mathbf{1}_\ell^T$

1120 *Proof of Lemma E.1.* Because \mathbf{Sub}_ℓ has linearly independent rows, the pseudo-inverse of it can be
1121 expressed as,

$$\mathbf{Sub}_\ell^\dagger = \mathbf{Sub}_\ell^T (\mathbf{Sub}_\ell \mathbf{Sub}_\ell^T)^{-1}$$

1122 From lemma D.1 we get,

$$\begin{aligned} \mathbf{Sub}_\ell^\dagger \mathbf{Sub}_\ell &= \mathbf{Sub}_\ell^T (\mathbf{Sub}_\ell \mathbf{Sub}_\ell^T)^{-1} \mathbf{Sub}_\ell \\ &= \mathcal{I}_\ell - \frac{1}{\ell} \mathbf{1}_\ell \mathbf{1}_\ell^T \end{aligned}$$

1123 Therefore,

$$\begin{aligned} \mathbf{Sub}_\ell^\dagger (\mathbf{Sub}_\ell \mathbf{Sub}_\ell^T) \mathbf{Sub}_\ell^{\dagger T} &= (\mathbf{Sub}_\ell^\dagger \mathbf{Sub}_\ell) (\mathbf{Sub}_\ell^\dagger \mathbf{Sub}_\ell)^T \\ &= (\mathcal{I}_\ell - \frac{1}{\ell} \mathbf{1}_\ell \mathbf{1}_\ell^T) (\mathcal{I}_\ell - \frac{1}{\ell} \mathbf{1}_\ell \mathbf{1}_\ell^T) \\ &= \mathcal{I}_\ell - \frac{1}{\ell} \mathbf{1}_\ell \mathbf{1}_\ell^T - \frac{1}{\ell} \mathbf{1}_\ell \mathbf{1}_\ell^T + \frac{1}{\ell^2} \mathbf{1}_\ell (\ell) \mathbf{1}_\ell^T \\ &= \mathcal{I}_\ell - \frac{1}{\ell} \mathbf{1}_\ell \mathbf{1}_\ell^T \end{aligned}$$

1124 \square

1125 THEOREM E.2. Let \mathbf{A} be a set of attributes and let $\mathbf{Q}_{\mathbf{A}}$ be the matrix representation of the marginal
 1126 on \mathbf{A} . Given the matrices $\mathbf{R}_{\mathbf{A}'}$, for all $\mathbf{A}' \in \text{closure}(\mathbf{A})$, we have $\mathbf{Q}_{\mathbf{A}} = \sum_{\mathbf{A}' \in \text{closure}(\mathbf{A})} \mathbf{Q}_{\mathbf{A}} \mathbf{R}_{\mathbf{A}'}^{\dagger} \mathbf{R}_{\mathbf{A}'}$.

Proof of Theorem E.2.

$$\mathbf{Q}_{\mathbf{A}} = \bigotimes_{i=1}^{n_a} \mathbf{K}_i \quad \text{where, for each } i, \mathbf{K}_i = \begin{cases} \mathcal{I}_{|Att_i|} & \text{if } Att_i \in \mathbf{A} \\ \mathbf{1}_{|Att_i|}^T & \text{if } Att_i \notin \mathbf{A} \end{cases}$$

$$\mathbf{R}_{\mathbf{A}'} = \bigotimes_{i=1}^{n_a} \mathbf{V}_i \quad \text{where, for each } i, \mathbf{V}_i = \begin{cases} \mathbf{Sub}_{|Att_i|} & \text{if } Att_i \in \mathbf{A}' \\ \mathbf{1}_{|Att_i|}^T & \text{if } Att_i \notin \mathbf{A}' \end{cases}$$

It is straightforward to verify that the following is a right inverse (and hence pseudo-inverse) of $\mathbf{R}_{\mathbf{A}'}$

$$\mathbf{R}_{\mathbf{A}'}^{\dagger} = \bigotimes_{i=1}^{n_a} \mathbf{V}_i^{\dagger} \quad \text{where, for each } i, \mathbf{V}_i^{\dagger} = \begin{cases} \mathbf{Sub}_{|Att_i|}^{\dagger} & \text{if } Att_i \in \mathbf{A}' \\ \frac{1}{|Att_i|} \mathbf{1}_{|Att_i|} & \text{if } Att_i \notin \mathbf{A}' \end{cases}$$

$$\mathbf{Q}_{\mathbf{A}} \mathbf{R}_{\mathbf{A}'}^{\dagger} \mathbf{R}_{\mathbf{A}'} = \bigotimes_{i=1}^{n_a} \mathbf{K}_i \mathbf{V}_i^{\dagger} \mathbf{V}_i \quad \text{where, for each } i, \mathbf{K}_i \mathbf{V}_i^{\dagger} \mathbf{V}_i = \begin{cases} \mathbf{Sub}_{|Att_i|}^{\dagger} \mathbf{Sub}_{|Att_i|} & \text{if } Att_i \in \mathbf{A}' \\ \frac{1}{|Att_i|} \mathbf{1}_{|Att_i|} \mathbf{1}_{|Att_i|}^T & \text{if } Att_i \in \mathbf{A}/\mathbf{A}' \\ \mathbf{1}_{|Att_i|}^T & \text{if } Att_i \notin \mathbf{A} \end{cases}$$

1127 Because $\mathbf{Sub}_{|Att_i|}$ has linearly independent rows, the pseudo-inverse of it can be expressed as,

$$\mathbf{Sub}_{|Att_i|}^{\dagger} = \mathbf{Sub}_{|Att_i|}^T (\mathbf{Sub}_{|Att_i|} \mathbf{Sub}_{|Att_i|}^T)^{-1}$$

1128 From lemma D.1 we get,

$$\begin{aligned} \mathbf{Sub}_{|Att_i|}^{\dagger} \mathbf{Sub}_{|Att_i|} &= \mathbf{Sub}_{|Att_i|}^T (\mathbf{Sub}_{|Att_i|} \mathbf{Sub}_{|Att_i|}^T)^{-1} \mathbf{Sub}_{|Att_i|} \\ &= \mathcal{I}_{|Att_i|} - \frac{1}{|Att_i|} \mathbf{1}_{|Att_i|} \mathbf{1}_{|Att_i|}^T \end{aligned}$$

1129 Therefore,

$$\mathbf{Q}_{\mathbf{A}} \mathbf{R}_{\mathbf{A}'}^{\dagger} \mathbf{R}_{\mathbf{A}'} = \bigotimes_{i=1}^{n_a} \mathbf{T}_i \quad \text{where, for each } i, \mathbf{T}_i = \begin{cases} \mathcal{I}_{|Att_i|} - \frac{1}{|Att_i|} \mathbf{1}_{|Att_i|} \mathbf{1}_{|Att_i|}^T & \text{if } Att_i \in \mathbf{A}' \\ \frac{1}{|Att_i|} \mathbf{1}_{|Att_i|} \mathbf{1}_{|Att_i|}^T & \text{if } Att_i \in \mathbf{A}/\mathbf{A}' \\ \mathbf{1}_{|Att_i|}^T & \text{if } Att_i \notin \mathbf{A} \end{cases}$$

1130 Without loss of generality (and to simplify notation), assume $\mathbf{A} = \{Att_1, \dots, Att_{\ell}\}$ consists of the
 1131 first ℓ attributes,

$$\begin{aligned} \mathbf{Q}_{\mathbf{A}} &= \left(\bigotimes_{i=1}^{\ell} \mathcal{I}_{|Att_i|} \right) \otimes \left(\bigotimes_{i=\ell+1}^{n_a} \mathbf{1}_{|Att_i|}^T \right) \\ \sum_{\mathbf{A}' \in \text{closure}(\mathbf{A})} \mathbf{Q}_{\mathbf{A}} \mathbf{R}_{\mathbf{A}'}^{\dagger} \mathbf{R}_{\mathbf{A}'} &= \sum_{\mathbf{A}' \in \text{closure}(\mathbf{A})} \left(\bigotimes_{i=1}^{n_a} \mathbf{T}_i \right) \\ &= \sum_{\mathbf{A}' \in \text{closure}(\mathbf{A})} \left(\left(\bigotimes_{i=1}^{\ell} \mathbf{T}_i \right) \otimes \left(\bigotimes_{i=\ell+1}^{n_a} \mathbf{1}_{|Att_i|}^T \right) \right) \\ &= \left(\sum_{\mathbf{A}' \in \text{closure}(\mathbf{A})} \left(\bigotimes_{i=1}^{\ell} \mathbf{T}_i \right) \right) \otimes \left(\bigotimes_{i=\ell+1}^{n_a} \mathbf{1}_{|Att_i|}^T \right) \\ &\quad \text{where, for each } i \leq \ell, \mathbf{T}_i = \begin{cases} \mathcal{I}_{|Att_i|} - \frac{1}{|Att_i|} \mathbf{1}_{|Att_i|} \mathbf{1}_{|Att_i|}^T & \text{if } Att_i \in \mathbf{A}' \\ \frac{1}{|Att_i|} \mathbf{1}_{|Att_i|} \mathbf{1}_{|Att_i|}^T & \text{if } Att_i \in \mathbf{A}/\mathbf{A}' \end{cases} \end{aligned}$$

1132 Because of the distributive property of the Kronecker product,

$$\begin{aligned} \bigotimes_{i=1}^{\ell} \mathcal{I}_{|Att_i|} &= \bigotimes_{i=1}^{\ell} \left(\left(\mathcal{I}_{|Att_i|} - \frac{1}{|Att_i|} \mathbf{1}_{|Att_i|} \mathbf{1}_{|Att_i|}^T \right) + \frac{1}{|Att_i|} \mathbf{1}_{|Att_i|} \mathbf{1}_{|Att_i|}^T \right) \\ &= \sum_{\mathbf{A}' \in \text{closure}(\mathbf{A})} \left(\bigotimes_{i=1}^{\ell} \mathbf{T}_i \right) \end{aligned}$$

1133 Therefore, combining everything together,

$$\begin{aligned} \sum_{\mathbf{A}' \in \text{closure}(\mathbf{A})} \mathbf{Q}_{\mathbf{A}} \mathbf{R}_{\mathbf{A}'}^{\dagger} \mathbf{R}_{\mathbf{A}'} &= \left(\sum_{\mathbf{A}' \in \text{closure}(\mathbf{A})} \left(\bigotimes_{i=1}^{\ell} \mathbf{T}_i \right) \right) \otimes \left(\bigotimes_{i=\ell+1}^{n_a} \mathbf{1}_{|Att_i|}^T \right) \\ &= \left(\bigotimes_{i=1}^{\ell} \mathcal{I}_{|Att_i|} \right) \otimes \left(\bigotimes_{i=\ell+1}^{n_a} \mathbf{1}_{|Att_i|}^T \right) \\ &= \mathbf{Q}_{\mathbf{A}} \end{aligned}$$

1134

□

1135 **THEOREM 4.7.** *Given a marginal workload $Wkload$ and positive numbers $\sigma_{\mathbf{A}}^2$ for each $\mathbf{A} \in$*
 1136 *closure($Wkload$), let \mathcal{M} be the mechanism that outputs $\{\mathcal{M}_{\mathbf{A}}(\mathbf{x}; \sigma_{\mathbf{A}}^2) : \mathbf{A} \in \text{closure}(Wkload)\}$*
 1137 *and let $\{\omega_{\mathbf{A}} : \mathbf{A} \in \text{closure}(Wkload)\}$ denote the privacy-preserving noisy answers (e.g.,*
 1138 *$\omega_{\mathbf{A}} = \mathcal{M}_{\mathbf{A}}(\mathbf{x}, \sigma^2)$). Then for any marginal on an attribute set $\mathbf{A} \in \text{closure}(Wkload)$, Algo-*
 1139 *algorithm 2 returns the unique linear unbiased estimate of $\mathbf{Q}_{\mathbf{A}} \mathbf{x}$ (i.e., answers to the marginal query)*
 1140 *that can be computed from the noisy differentially private answers.*

1141 *The variances $\text{Var}(\mathbf{A}; \mathcal{M})$ of all the noisy cell counts of the marginal on \mathbf{A} is the vector*
 1142 *whose components are all equal to $\sum_{\mathbf{A}' \subseteq \mathbf{A}} \left(\sigma_{\mathbf{A}'}^2 \prod_{Att_i \in \mathbf{A}'} \frac{|Att_i|-1}{|Att_i|} * \prod_{Att_j \in (\mathbf{A}/\mathbf{A}')} \frac{1}{|Att_j|^2} \right)$.*
 1143 *The covariance between any two noisy answers of the marginal on \mathbf{A} is*
 1144 $\sum_{\mathbf{A}' \subseteq \mathbf{A}} \left(\sigma_{\mathbf{A}'}^2 \prod_{Att_i \in \mathbf{A}'} \frac{-1}{|Att_i|} * \prod_{Att_j \in (\mathbf{A}/\mathbf{A}')} \frac{1}{|Att_j|^2} \right)$.

1145 *Proof of Theorem 4.7.* We first verify the correctness and uniqueness of the reconstruction in
 1146 Algorithm 2. Uniqueness follows from the fact that the rows from all the matrices $\mathbf{R}_{\mathbf{A}}$ (for
 1147 $\mathbf{A} \in \text{closure}(Wkload)$) are linearly independent.

1148 Consider Line 3 from Algorithm 2. It uses a \mathbf{U} matrix that depends on both the attributes \mathbf{A}
 1149 of the marginal one wants to compute and a subset \mathbf{A}' of it. So, for notational dependence, we
 1150 write it as $\mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'}$. It is straightforward to verify that $\mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'} = \mathbf{Q}_{\mathbf{A}} \mathbf{R}_{\mathbf{A}'}^{\dagger}$. From Theorem E.2,
 1151 $\mathbf{Q}_{\mathbf{A}} \mathbf{x} = \sum_{\mathbf{A}' \subseteq \mathbf{A}} \mathbf{Q}_{\mathbf{A}} \mathbf{R}_{\mathbf{A}'}^{\dagger} \mathbf{R}_{\mathbf{A}'} \mathbf{x} = \sum_{\mathbf{A}' \subseteq \mathbf{A}} \mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'} \mathbf{R}_{\mathbf{A}'} \mathbf{x}$, and so Algorithm 2 is correct because
 1152 each $\omega_{\mathbf{A}'}$ is an unbiased noisy version of $\mathbf{R}_{\mathbf{A}'} \mathbf{x}$.

1153 Having established that the \mathbf{q} returned by Line 5 in Algorithm 2 is an unbiased estimate of the
 1154 marginal query answer $\mathbf{Q}_{\mathbf{A}} \mathbf{x}$, the next step is to compute the covariance matrix $E[\mathbf{q}\mathbf{q}^T]$.

$$\begin{aligned} E[\mathbf{q}\mathbf{q}^T] &= E \left[\sum_{\mathbf{A}' \subseteq \mathbf{A}} \mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'} (\omega_{\mathbf{A}'} \omega_{\mathbf{A}'}^T) \mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'}^T \right] \\ &= \sum_{\mathbf{A}' \subseteq \mathbf{A}} \mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'} (\sigma_{\mathbf{A}'}^2 \Sigma_{\mathbf{A}'}) \mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'}^T \end{aligned}$$

1155 Without loss of generality (and to simplify notation), assume $\mathbf{A} = \{Att_1, \dots, Att_{\ell}\}$ consists of
 1156 the first ℓ attributes, $\mathbf{A}' = \{Att_1, \dots, Att_t\}$ consists of the first $t \leq \ell$ attributes, then $\mathbf{A}/\mathbf{A}' =$
 1157 $\{Att_{t+1}, \dots, Att_{\ell}\}$.

1158 By definition, $\text{Var}(A; \mathcal{M})$ is the diagonal of $E[\mathbf{q}\mathbf{q}^T] = \sum_{\mathbf{A}' \in \text{closure}(\mathbf{A})} \sigma_{\mathbf{A}'}^2 \mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'} \Sigma_{\mathbf{A}'} \mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'}^T$.
 1159 Thus we can write:

$$\begin{aligned}
 \mathbf{Q}_{\mathbf{A}} &= \left(\bigotimes_{i=1}^t \mathcal{I}_{|Att_i|} \right) \otimes \left(\bigotimes_{j=t+1}^{\ell} \mathcal{I}_{|Att_j|} \right) \otimes \left(\bigotimes_{k=\ell+1}^{n_a} \mathbf{1}_{|Att_k|}^T \right) \\
 \mathbf{R}_{\mathbf{A}'} &= \left(\bigotimes_{i=1}^t \text{Sub}_{|Att_i|} \right) \otimes \left(\bigotimes_{j=t+1}^{\ell} \mathbf{1}_{|Att_j|}^T \right) \otimes \left(\bigotimes_{k=\ell+1}^{n_a} \mathbf{1}_{|Att_k|}^T \right) \\
 \mathbf{R}_{\mathbf{A}'}^{\dagger} &= \left(\bigotimes_{i=1}^t \text{Sub}_{|Att_i|}^{\dagger} \right) \otimes \left(\bigotimes_{j=t+1}^{\ell} \frac{1}{|Att_j|} \mathbf{1}_{|Att_j|} \right) \otimes \left(\bigotimes_{k=\ell+1}^{n_a} \frac{1}{|Att_k|} \mathbf{1}_{|Att_k|} \right) \\
 \mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'} &= \mathbf{Q}_{\mathbf{A}} \mathbf{R}_{\mathbf{A}'}^{\dagger} = \left(\bigotimes_{i=1}^t \text{Sub}_{|Att_i|}^{\dagger} \right) \otimes \left(\bigotimes_{j=t+1}^{\ell} \frac{1}{|Att_j|} \mathbf{1}_{|Att_j|} \right) \otimes \left(\bigotimes_{k=\ell+1}^{n_a} [1] \right) \\
 \mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'}^T &= \left(\bigotimes_{i=1}^t \text{Sub}_{|Att_i|}^{\dagger T} \right) \otimes \left(\bigotimes_{j=t+1}^{\ell} \frac{1}{|Att_j|} \mathbf{1}_{|Att_j|}^T \right) \otimes \left(\bigotimes_{k=\ell+1}^{n_a} [1] \right) \\
 \Sigma_{\mathbf{A}'} &= \left(\bigotimes_{i=1}^t \text{Sub}_{|Att_i|} \text{Sub}_{|Att_i|}^T \right) \otimes \left(\bigotimes_{j=t+1}^{\ell} [1] \right) \otimes \left(\bigotimes_{k=\ell+1}^{n_a} [1] \right) \\
 \mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'} \Sigma_{\mathbf{A}'} \mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'}^T &= \left(\bigotimes_{i=1}^t \text{Sub}_{|Att_i|}^{\dagger} \text{Sub}_{|Att_i|} \text{Sub}_{|Att_i|}^T \text{Sub}_{|Att_i|}^{\dagger T} \right) \\
 &\quad \otimes \left(\bigotimes_{j=t+1}^{\ell} \frac{1}{|Att_j|^2} \mathbf{1}_{|Att_j|} [1] \mathbf{1}_{|Att_j|}^T \right) \otimes \left(\bigotimes_{k=\ell+1}^{n_a} [1] \right) \tag{10}
 \end{aligned}$$

1160 Now, by Lemma E.1,

$$\text{Sub}_{|Att_i|}^{\dagger} \text{Sub}_{|Att_i|} \text{Sub}_{|Att_i|}^T \text{Sub}_{|Att_i|}^{\dagger T} = \mathcal{I}_{|Att_i|} - \frac{1}{|Att_i|} \mathbf{1}_{|Att_i|} \mathbf{1}_{|Att_i|}^T \tag{11}$$

1161 So the diagonals of $\mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'} \Sigma_{\mathbf{A}'} \mathbf{U}_{\mathbf{A} \leftarrow \mathbf{A}'}^T$ can be computed by multiplying $\frac{|Att_i|-1}{|Att_i|}$ for each $Att_i \in$
 1162 \mathbf{A}' and $1/|Att_j|$ for each $Att_j \in \mathbf{A} \setminus \mathbf{A}'$. Meanwhile, the off diagonals are all the same and can be
 1163 computed by multiplying $\frac{-1}{|Att_i|}$ for each $Att_i \in \mathbf{A}'$ and $\frac{1}{|Att_j|^2}$ for each $Att_j \in \mathbf{A} \setminus \mathbf{A}'$.

1164 Computing the variance and covariance of the marginal query answer is therefore the summation of
 1165 these quantities for all $\mathbf{A}' \subseteq \mathbf{A}$ and is what the theorem states.

1166 □

1167 F Computational Complexity Proofs

1168 **THEOREM 4.8.** *Let n_a be the total number of attributes. Let $\#\text{cells}(\mathbf{A})$ denote the number of cells in*
 1169 *the marginal on attribute set \mathbf{A} . Then:*

- 1170 1. *Expressing the privacy cost of the optimal mechanism \mathcal{M}^* as a linear combination of the $1/\sigma_{\mathbf{A}}^2$*
 1171 *values takes $O(\sum_{\mathbf{A} \in \text{Wkload}} \#\text{cells}(\mathbf{A}))$ total time.*
- 1172 2. *Expressing all of the $\text{Var}(\mathbf{A}; \mathcal{M}^*)$, for $\mathbf{A} \in \text{Wkload}$, as a linear combinations of the $\sigma_{\mathbf{A}}^2$ values*
 1173 *can be done in $O(\sum_{\mathbf{A} \in \text{Wkload}} \#\text{cells}(\mathbf{A}))$ total time.*
- 1174 3. *Computing all the noisy outputs of the optimal mechanism (i.e., $\mathcal{M}_{\mathbf{A}}(\mathbf{x}; \sigma_{\mathbf{A}}^2)$ for $\mathbf{A} \in$*
 1175 *$\text{closure}(\text{Wkload})$) takes $O(n_a \sum_{\mathbf{A} \in \text{Wkload}} \prod_{Att_i \in \mathbf{A}} (|Att_i| + 1))$ total time after the true an-*
 1176 *swers have been precomputed (Line 1 in Algorithm 1). Note that the total number of cells on*
 1177 *marginals in Wkload is $O(\sum_{\mathbf{A} \in \text{Wkload}} \prod_{Att_i \in \mathbf{A}} |Att_i|)$.*

- 1178 4. Reconstructing marginals for all $\mathbf{A} \in Wkload$ takes $O(\sum_{\mathbf{A} \in Wkload} |\mathbf{A}| \#cells(\mathbf{A})^2)$ total time.
 1179 5. Computing the variance of the cells for all of the marginals for $\mathbf{A} \in Wkload$ can be done in
 1180 $O(\sum_{\mathbf{A} \in Wkload} \#cells(\mathbf{A}))$ total time.

1181 *Proof of Theorem 4.8.* First we establish that $|\text{closure}(Wkload)| \leq \sum_{\mathbf{A} \in Wkload} \#cells(\mathbf{A})$. Given
 1182 an set $\mathbf{A} \in Wkload$, we note that it has $2^{|\mathbf{A}|}$ subsets, so that $|\text{closure}(\mathbf{A})| = 2^{|\mathbf{A}|}$. However,
 1183 $\#cells(\mathbf{A})$ is at least $2^{|\mathbf{A}|}$ (because each attribute has at least 2 attribute values). We also note that
 1184 $\text{closure}(Wkload) = \bigcup_{\mathbf{A} \in Wkload} \text{closure}(\mathbf{A})$. Hence

$$|\text{closure}(Wkload)| \leq \sum_{\mathbf{A} \in Wkload} |\text{closure}(\mathbf{A})| = \sum_{\mathbf{A} \in Wkload} \#cells(\mathbf{A})$$

1185 To analyze the time complexity of symbolically representing the privacy cost, as a linear combi-
 1186 nation of the $1/\sigma_{\mathbf{A}}^2$ values (for all $\mathbf{A} \in \text{closure}(Wkload)$) we note that the coefficient of $1/\sigma_{\mathbf{A}}^2$
 1187 is $\prod_{Att_i \in \mathbf{A}} \frac{|Att_i|-1}{|Att_i|}$. Thus computing the coefficient $1/\sigma_{\emptyset}^2$ takes $O(1)$ time. Then, computing the
 1188 coefficient of $1/\sigma_{\{Att_i\}}^2$ can be computed from the coefficient of $1/\sigma_{\emptyset}^2$ in $O(1)$ additional time. Thus,
 1189 we if go level by level, first computing the coefficients of $1/\sigma_{\mathbf{A}}^2$ with $|\mathbf{A}| = 1$ then for $|\mathbf{A}| = 2$, etc.
 1190 then computing the coefficient for each new \mathbf{A} takes incremental $O(1)$ time. Thus the overall time is
 1191 $O(|\text{closure}(Wkload)|)$ and therefore is $O(\sum_{\mathbf{A} \in Wkload} \#cells(\mathbf{A}))$.

1192 Let $n_{cells} = \sum_{\mathbf{A} \in Wkload} \#cells(\mathbf{A})$ To express the variance symbolically as a linear function of
 1193 the $\sigma_{\mathbf{A}}^2$ values via Theorem 4.7, we note from the previous part that computing $\prod_{Att_i \in \mathbf{A}'} \frac{|Att_i|-1}{|Att_i|}$ for
 1194 all $\mathbf{A}' \in \text{closure}(Wkload)$ can be done in total $O(n_{cells})$ time. Similarly, computing $\prod_{Att_i \in \mathbf{A}'} \frac{1}{|Att_i|^2}$
 1195 for all $\mathbf{A}' \in \text{closure}(Wkload)$ also take total $O(n_{cells})$ time. Once this is pre-computed, then
 1196 for any $\mathbf{A}' \subseteq \mathbf{A} \in \text{closure}(Wkload)$, the product $\prod_{Att_i \in \mathbf{A}'} \frac{|Att_i|-1}{|Att_i|} * \prod_{Att_j \in (\mathbf{A}/\mathbf{A}')} \frac{1}{|Att_j|^2}$
 1197 can be computed in $O(1)$ time since $\mathbf{A} \setminus \mathbf{A}' \in \text{closure}(Wkload)$. Now, $Var(\mathbf{A}; \mathcal{M}^*) =$
 1198 $\sum_{\mathbf{A}' \subseteq \mathbf{A}} \sigma_{\mathbf{A}'}^2 \prod_{Att_i \in \mathbf{A}'} \frac{|Att_i|-1}{|Att_i|} * \prod_{Att_j \in (\mathbf{A}/\mathbf{A}')} \frac{1}{|Att_j|^2}$. This is a linear combination of $2^{|\mathbf{A}|}$ terms
 1199 (one term for each variable $\sigma_{\mathbf{A}'}$, for $\mathbf{A}' \subseteq \mathbf{A}$). Each term is computed in $O(1)$ time after the pre-
 1200 computation phase. Thus the symbolic representation of $Var(\mathbf{A}; \mathcal{M}^*)$ takes $O(2^{|\mathbf{A}|})$ time (which
 1201 is at most the number of cells in the marginal on \mathbf{A}) time after precomputation. Thus computing
 1202 $Var(\mathbf{A}; \mathcal{M}^*)$ for all $\mathbf{A} \in Wkload$ can be done in total $O(n_{cells})$ time after precomputation, but
 1203 precomputation also takes $O(n_{cells})$ time. Thus the overall total time is $O(n_{cells})$.

1204 We next analyze the time it takes to generate noisy answers once the true answers have been
 1205 precomputed (Line 1 in Algorithm 1). This involves (1) computing the product $\mathbf{H}\mathbf{v}$ in the algorithm,
 1206 (2) generating one Gaussian random variable for each column of \mathbf{H} and (3) computing $\mathbf{H}\mathbf{z}$. Now, the
 1207 first and third steps take the same amount of time. The second step generates one Gaussian for each
 1208 row of \mathbf{H} and hence, for each $\mathcal{M}_{\mathbf{A}}$ takes time $\prod_{Att_i \in \mathbf{A}} (|Att_i| - 1)$.

1209 For the first step, the fast kronecker-product multiplication algorithm (Algorithm 1 of [38]) has the
 1210 following complexity. Given a kronecker product of ℓ matrices of sizes $(m_1 - 1) \times m_1, \dots, (m_{\ell} - 1) \times$
 1211 m_{ℓ} and a vector with $m_1 \times \dots \times m_{\ell}$ components, their algorithm has ℓ iterations. In iteration i , the i^{th}
 1212 matrix (with size $m_{i-1} \times m_i$) is multiplied by a matrix with shape $(m_i, \prod_{j=1}^{i-1} m_j * \prod_{j=i+1}^{\ell} (m_j - 1))$.
 1213 In our case, each m_i is a subtraction matrix with two nonzero elements in each row. Thus, in each
 1214 iteration, the product makes $2 \prod_{j=1}^{i-1} m_j * \prod_{j=i}^{\ell} (m_j - 1)$ scalar multiplication operations. There are
 1215 ℓ iterations, so the multiplication algorithm uses $O(\ell \prod_{i=1}^{\ell} m_i)$ multiplications.

1216 Now, to run algorithm $\mathcal{M}_{\mathbf{A}}$, the number of kron products ℓ is $|\mathbf{A}|$ and each m_i is $|Att_i|$ for $Att_i \in \mathbf{A}$.
 1217 Hence the running time of $\mathcal{M}_{\mathbf{A}}$ is $O(|\mathbf{A}| \prod_{Att_i \in \mathbf{A}} |Att_i|)$ which is at most $|\mathbf{A}|$ times the number
 1218 of cells in the marginal on \mathbf{A} . Note that the constant in the big-O notation is bounded across all \mathbf{A} .

1219 Next, when adding up the complexity across all $\mathbf{A}' \in \text{closure}(\mathbf{A})$, we can replace $|\mathbf{A}'|$ with $|\mathbf{A}|$, and
 1220 then the summation looks like the product $\prod_{Att_i \in \mathbf{A}} (|Att_i| + 1)$ when this product is expanded. Hence
 1221 the time to run all $\mathbf{Q}_{\mathbf{A}'}$ for all $\mathbf{A}' \in \text{closure}(\mathbf{A})$ is $O(|\mathbf{A}| \prod_{Att_i \in \mathbf{A}} (|Att_i| + 1))$. Adding up over all
 1222 $\mathbf{A} \in Wkload$ gets the results.

1223 Next we consider the reconstruction phase. Using the same analysis of the fast kron-product
 1224 vector multiplication, we see that in each iteration of Algorithm 2, there is a kron product vector
 1225 multiplication. Using similar reasoning as for the previous item, each such multiplication takes
 1226 $O(|\mathbf{A}| \prod_{Att_i \in \mathbf{A}} |Att_i|) = O(|\mathbf{A}| \#cells(\mathbf{A}))$ time. The number of iterations in the algorithm is
 1227 $2^{|\mathbf{A}|} \leq \#cells(\mathbf{A})$. Thus the overall runtime is $O(\sum_{\mathbf{A} \in Wkload} |\mathbf{A}| \#cells(\mathbf{A})^2)$.

1228 Finally, the variance computation is no harder than expressing the $Var(\mathbf{A}; \mathcal{M}^*)$ as linear combina-
 1229 tions of the optimization variables and we have shown this to be $O(n_{cells})$. \square

1230 G Closed Form Solution to the Weighted Sum of Variances Loss

1231 By Theorem 4.5, the privacy cost is a linear combination of the $1/\sigma_{\mathbf{A}}^2$ values. By Theorem 4.7, each
 1232 reconstructed marginal's cell variances are a linear combination of the $\sigma_{\mathbf{A}}^2$ values. Thus, minimizing
 1233 the weighted sum of reconstructed marginal variances subject to the privacy cost being $\leq c$ can be
 1234 formulated as a problem of the following type:

$$\begin{aligned} \sigma_{\mathbf{A}}^2 : \mathbf{A} \in \text{closure}(Wkload) \quad & \arg \min & \sum_{\mathbf{A} \in \text{closure}(Wkload)} v_{\mathbf{A}} \sigma_{\mathbf{A}}^2 & (12) \\ \text{s.t.} & & \sum_{\mathbf{A} \in \text{closure}(Wkload)} \frac{p_{\mathbf{A}}}{\sigma_{\mathbf{A}}^2} \leq c & \end{aligned}$$

1235 where the $v_{\mathbf{A}}$ are the linear coefficients of the $\sigma_{\mathbf{A}}^2$ and the $p_{\mathbf{A}}$ are the linear coefficients of the $1/\sigma_{\mathbf{A}}^2$
 1236 in the privacy cost. The closed form solution is given by the following lemma.

1237 LEMMA G.1. *Given the optimization problem in Equation 12 The optimal objective function value is*
 1238 $T = (\sum_{\mathbf{A}} \sqrt{v_{\mathbf{A}} p_{\mathbf{A}}})^2 / c$, *the optimal value of each noise scale parameter is $\sigma_{\mathbf{A}}^2 = \sqrt{T p_{\mathbf{A}} / (c v_{\mathbf{A}})}$.*

1239 *Proof.* Clearly, for the optimal solution, the inequality constraint must be tight (i.e., $= c$) because if
 1240 it is not tight, we can lower variance while increasing privacy cost by dividing each $\sigma_{\mathbf{A}}^2$ by a number
 1241 > 1 . Thus we just need to solve the problem subject to $\sum_{\mathbf{A}} p_{\mathbf{A}} / \sigma_{\mathbf{A}}^2 = c$.

1242 From Cauchy-Schwarz inequality,

$$\sum_{\mathbf{A}} v_{\mathbf{A}} \sigma_{\mathbf{A}}^2 = \left(\sum_{\mathbf{A}} v_{\mathbf{A}} \sigma_{\mathbf{A}}^2 \right) \left(\sum_{\mathbf{A}} \frac{p_{\mathbf{A}}}{\sigma_{\mathbf{A}}^2} \right) / c \geq \left(\sum_{\mathbf{A}} \sqrt{v_{\mathbf{A}} p_{\mathbf{A}}} \right)^2 / c = T$$

1243 Equality holds when $\frac{v_{\mathbf{A}}}{p_{\mathbf{A}}} \sigma_{\mathbf{A}}^4 = t$ for all \mathbf{A} (for some constant t). Since $c = \sum_{\mathbf{A}} \frac{p_{\mathbf{A}}}{\sigma_{\mathbf{A}}^2} =$
 1244 $\sum_{\mathbf{A}} \sqrt{v_{\mathbf{A}} p_{\mathbf{A}}} / t$, then we must have $t = T/c$. Plugging this into the definition of t , we get
 1245 $\sigma_{\mathbf{A}}^2 = \sqrt{T p_{\mathbf{A}} / (c v_{\mathbf{A}})}$. \square

1246 Thus, if the loss function is the weighted sum of variances, ResidualPlanner does not need any
 1247 optimization steps. The selection of the noise scales and the reconstruction phase are direct algorithms.

1248 H Additional Experiments

1249 In this section, we present additional experiments. Following [37], the experiments use the following
 1250 type of workloads:

- 1251 • All k -way marginals.

- 1252 • All ≤ 3 -way marginals. This includes all 0-way marginal (the total sum), all 1-way
1253 marginals, all 2-way marginals, and all 3-way marginals.
- 1254 • Small marginals. This includes any k -way marginal that has at most 5000 cells.

1255 We also use these metrics:

- 1256 • RMSE: The total variance is the sum of the variances of the reconstructed cells in each
1257 marginal in the workload. Root Mean Squared Error is obtained by taking the total variance,
1258 dividing by the total number of cells in the workload marginals, then taking the square root.
1259 The SVD Bound (SVDB for short) [31] provides a theoretical lower bound on RMSE for any
1260 matrix mechanism. For marginals, the SVDB is tight, but its computation is not scalable.
- 1261 • MaxVar: compute the variance of each reconstructed cell for each marginal in the workload,
1262 then take the maximum of these.
- 1263 • Running time (in seconds) of the different stages of the algorithms (select and reconstruct).

1264 Unless otherwise stated, ResidualPlanner uses the open-source ECOS optimizer [14] for solving the
1265 optimization problem it generates for the select step.

1266 For all experiments, we require all mechanisms to have privacy cost $pcost(\mathcal{M}) = 1$. By definition
1267 2.3, \mathcal{M} satisfies ρ -zCDP with $\rho = 1/2$ [46] and satisfies μ -Gaussian DP with $\mu = 1$ [15, 46].

1268 Each experiment is repeated 5 times, we report the mean value of these 5 results and a confidence
1269 interval consisting of ± 2 standard deviations. This is most useful for running time, as the variance
1270 loss metrics have negligible variance across all algorithms.

1271 H.1 Scalability

1272 In this section, we study the scalability of ResidualPlanner. This is done using the Synth- n^d dataset,
1273 where d is the number of attributes and n is the domain size of each attribute. We use all ≤ 3 -
1274 way marginals as a fixed workload and vary n or d to get the computation time for HDMM and
1275 ResidualPlanner.

1276 H.1.1 Varying Attribute Domain Size n in the Selection Step.

1277 This experiment considers what happens when the attribute domain size n get larger. We fix the
1278 number of attributes $d = 5$ and vary the domain size n for each attribute, where n ranges from 2 to
1279 1024. We evaluate the running time and accuracy of the selection step

1280 Table 10 shows the running time for the selection step of HDMM and ResidualPlanner. The RMSE on
1281 the workload that the selection step guarantees is also measured. Both HDMM and ResidualPlanner
1282 have no trouble here. HDMM is nearly optimal in RMSE and ResidualPlanner is optimal, as shown
1283 by agreement with the SVD Bound. ResidualPlanner is faster, but both methods are fast in this
1284 experiment setting.

Table 10: Selection step on Synth- n^d dataset where $d = 5$ and n varies. The workload is all ≤ 3 -way marginals. Metrics are running time and RMSE.

n	$Time_{HDMM}$	$Time_{ResPlan}$	$RMSE_{HDMM}$	$RMSE_{ResPlan}$	SVDB
2	0.069 ± 0.018	0.001 ± 0.000	1.903	1.890	1.890
4	0.064 ± 0.006	0.001 ± 0.000	2.685	2.681	2.681
8	0.070 ± 0.021	0.001 ± 0.000	3.156	3.156	3.156
16	0.076 ± 0.020	0.001 ± 0.000	3.367	3.366	3.366
32	0.105 ± 0.020	0.001 ± 0.000	3.422	3.423	3.423
64	0.114 ± 0.033	0.001 ± 0.000	3.408	3.407	3.407
128	0.137 ± 0.048	0.001 ± 0.000	3.371	3.367	3.367
256	0.187 ± 0.050	0.001 ± 0.000	3.331	3.322	3.322
512	0.183 ± 0.020	0.001 ± 0.000	3.294	3.283	3.283
1024	0.353 ± 0.058	0.001 ± 0.000	3.328	3.251	3.251

1285 Table 11 shows the running time and Max Variance comparison for the selection step. HDMM can
 1286 only optimize for RMSE, not max variance, so this table shows that RMSE is not a good substitute
 when one needs to optimize for Max Variance.

Table 11: Selection step on Synth- n^d dataset where $d = 5$ and n varies. The workload is all \leq
 3-way marginals. Metrics are running time and Max Variance.

n	$Time_{HDMM}$	$Time_{ResPlan}$	$MaxVar_{HDMM}$	$MaxVar_{ResPlan}$
2	0.069 ± 0.018	0.008 ± 0.001	8.091	4.148
4	0.064 ± 0.006	0.008 ± 0.001	44.693	9.760
8	0.070 ± 0.021	0.008 ± 0.001	180.343	15.643
16	0.076 ± 0.020	0.008 ± 0.001	588.115	20.067
32	0.105 ± 0.020	0.008 ± 0.001	1649.341	22.811
64	0.114 ± 0.033	0.008 ± 0.001	5560.807	24.345
128	0.137 ± 0.048	0.008 ± 0.001	12229.480	25.157
256	0.187 ± 0.050	0.008 ± 0.001	8168.716	25.574
512	0.183 ± 0.020	0.008 ± 0.001	32159.958	25.786
1024	0.353 ± 0.058	0.008 ± 0.001	277825.955	25.893

1287

1288 H.1.2 Impact of varying the number of attributes in the Selection Step.

1289 Next, we fix the domain size of each attribute to be $n = 10$ and vary the number of attributes d ,
 1290 where d ranges from 2 to 200. This experiment can test some of the limits of ResidualPlanner. While
 1291 HDMM cannot perform selection when the number of attributes is 20 or larger, ResidualPlanner has
 1292 no trouble optimizing RMSE even for 200 attributes. However, optimizing for Max Variance is much
 1293 more difficult. ResidualPlanner can do this for $d = 100$ but the underlying optimization took more
 1294 than 1 hour for $d = 200$ and we killed the process.

1295 Table 12 shows the running time and RMSE comparison for the selection step. The running time of
 1296 HDMM increases sharply and it quickly runs out of memory. At the same point, the SVD Bound can
 1297 no longer be computed. Meanwhile, ResidualPlanner continues to run efficiently.

Table 12: Selection step on Synth- n^d dataset where $n = 10$ and d varies. The workload is all \leq
 3-way marginals. Metrics are running time and RMSE.

d	$Time_{HDMM}$	$Time_{ResPlan}$	$RMSE_{HDMM}$	$RMSE_{ResPlan}$	SVDB
2	0.013 ± 0.003	0.001 ± 0.0008	1.379	1.379	1.379
4	0.028 ± 0.007	0.002 ± 0.001	2.346	2.345	2.345
6	0.065 ± 0.012	0.002 ± 0.0008	4.278	4.275	4.275
8	0.167 ± 0.019	0.004 ± 0.001	6.726	6.638	6.638
10	0.639 ± 0.059	0.009 ± 0.001	9.629	9.348	9.348
12	4.702 ± 0.315	0.015 ± 0.001	12.904	12.359	12.359
14	46.054 ± 12.735	0.025 ± 0.002	16.506	15.642	15.642
15	201.485 ± 13.697	0.030 ± 0.017	18.421	17.378	17.378
20	Out of memory	0.079 ± 0.017	Out of memory	26.916	Out of memory
30	Out of memory	0.247 ± 0.019	Out of memory	49.713	Out of memory
50	Out of memory	1.207 ± 0.047	Out of memory	107.258	Out of memory
100	Out of memory	9.913 ± 0.246	Out of memory	303.216	Out of memory
200	Out of memory	80.120 ± 1.502	Out of memory	855.330	Out of memory

1298 Table 13 shows the running time and Max Variance comparison on the Selection step. Optimizing
 1299 for Max Variance is much harder for ResidualPlanner compared to RMSE and we killed the process
 1300 for $d = 200$. Meanwhile, HDMM is not able to run at $d = 20$ (we emphasize again, it optimizes for
 1301 RMSE even if one cares about Max Variance). There is an interesting phenomenon with HDMM that
 1302 takes place for d between 8 and 15. In this case, HDMM always produces a max variance of 1000.
 1303 This maximum is always achieved for the sum query (a zero-dimensional marginal) for the following
 1304 reason. For d between 8 and 15, HDMM decides to add noise to all 3-way marginals and nothing else
 1305 (even though the workload is all ≤ 3 marginals). The privacy loss budget is split equally among them.
 1306 Thus, each of the $\binom{d}{3}$ marginals it measures gets $N(0, \binom{d}{3})$ noise. The sum query gets reconstructed

1307 as follows. For any single noisy 3-way marginal, one can estimate the sum by adding up the cells in
 1308 the marginal. Since each cell has variance $\binom{d}{3}$ and there are $n^3 = 1,000$ cells, the sum estimate from
 1309 a single 3-way marginal has a variance of $1000\binom{d}{3}$. But one can obtain an independent estimate to
 1310 the sum query from each of the $\binom{d}{3}$ noisy 3-way marginals. By averaging these noisy estimates, one
 1311 can obtain an estimate of the sum query with variance 1,000.

Table 13: Selection step on Synth- n^d dataset where $n = 10$ and d varies. The workload is all \leq 3-way marginals. Metrics are running time and Max Variance.

d	$Time_{HDMM}$	$Time_{ResPlan}$	$MaxVar_{HDMM}$	$MaxVar_{ResPlan}$
2	0.013 ± 0.003	0.007 ± 0.001	13.745	3.306
4	0.028 ± 0.007	0.010 ± 0.005	132.620	10.480
6	0.065 ± 0.012	0.009 ± 0.001	461.132	26.904
8	0.167 ± 0.019	0.015 ± 0.003	1000.000	56.961
10	0.639 ± 0.059	0.018 ± 0.001	1000.000	105.031
12	4.702 ± 0.315	0.028 ± 0.001	1000.000	175.496
14	46.054 ± 12.735	0.041 ± 0.001	1000.000	272.738
15	201.485 ± 13.697	0.050 ± 0.001	1000.000	332.769
20	Out of memory	0.123 ± 0.023	Out of memory	768.941
30	Out of memory	0.461 ± 0.024	Out of memory	2540.440
50	Out of memory	4.011 ± 0.112	Out of memory	11597.037
100	Out of memory	121.224 ± 3.008	Out of memory	91960.917

1312 H.1.3 Scalability of the Reconstruction Step.

1313 We conduct similar experiments, but now we measure the time in the reconstruction step. To com-
 1314 plement the reconstruction scalability experiments from the main paper on the Synth- n^d synthetic
 1315 dataset, we first fix the number of attributes $d = 5$ and vary the domain size n for each attribute,
 1316 where n ranges from 2 to 512. The reconstruction time for ResidualPlanner does not depend on the
 1317 metric that the select step was optimized for. Again we compare with HDMM [38] and a version of
 1318 HDMM with improved reconstruction scalability called HDMM+PGM [38, 41] (the PGM settings
 1319 used 50 iterations of its Local-Inference estimator, as the default 1000 was too slow). Table 14 shows
 1320 the results. Again, at some point HDMM runs out of memory while ResidualPlanner runs efficiently.
 1321 HDMM runs out of memory because of choices it had made in the selection step. When $n = 128$ it
 1322 decided to measure a 5-way marginal, which is so large (requiring 128^5 space) that it caused HDMM
 and HDMM+PGM to have memory issues.

Table 14: Running time (in seconds) of the reconstruction step on Synth- n^d dataset where $d = 5$ and n varies. The workload is all \leq 3-way marginals.

n	HDMM	HDMM + PGM	ResPlan
2	0.005 ± 0.002	2.466 ± 0.278	0.008 ± 0.002
4	0.005 ± 0.000	1.894 ± 0.146	0.011 ± 0.008
8	0.008 ± 0.000	1.871 ± 0.122	0.011 ± 0.008
16	0.064 ± 0.036	1.936 ± 0.131	0.016 ± 0.001
32	1.924 ± 0.060	3.211 ± 0.220	0.045 ± 0.007
64	56.736 ± 1.460	12.574 ± 0.512	0.217 ± 0.021
128	Out of memory	Out of memory	1.244 ± 0.059
256	Out of memory	Out of memory	12.090 ± 0.504
512	Out of memory	Out of memory	166.045 ± 13.803

1323
 1324 We next fix $n = 3$ and vary d . Table 15 shows ResidualPlanner is clearly faster. Furthermore, HDMM
 1325 and HDMM+PGM are hampered by the failure of the selection step (when selection fails, there is
 1326 nothing to reconstruct). It is interesting to compare HDMM+PGM behavior when $n = 3$ in Table
 1327 15 with $n = 10$ in Table 2 from the main paper. Clearly HDMM+PGM is faster for $n = 10$ than
 1328 $n = 3$. This counterintuitive result can be explained by the complex workings of HDMM as follows.
 1329 When $n = 3$, the selection step in HDMM returns some 4-way marginals. But when $n = 10$, HDMM

1330 only returns ≤ 3 -way marginals. The 4-way marginals make the reconstruction step harder for both
 1331 HDMM and HDMM + PGM.

Table 15: **Time for Reconstruction Step in seconds** on Synth- n^d dataset. $n = 3$ and the number of attributes d varies. The workload consists of all marginals on ≤ 3 attributes each. Times are reported with ± 2 standard deviations. Reconstruction can only be performed if the select step completed.

d	HDMM	HDMM + PGM	ResidualPlanner
2	0.001 \pm 0.0001	0.256 \pm 0.030	0.005 \pm 0.002
6	0.009 \pm 0.001	3.293 \pm 0.253	0.020 \pm 0.004
10	0.334 \pm 0.010	51.568 \pm 3.391	0.086 \pm 0.004
12	3.882 \pm 0.101	180.708 \pm 5.437	0.153 \pm 0.002
14	55.856 \pm 0.361	314.252 \pm 3.991	0.280 \pm 0.072
15	231.283 \pm 0.554	713.526 \pm 4.957	0.307 \pm 0.005
20	Unavailable (select step failed)	Unavailable (select step failed)	0.758 \pm 0.023
30	Unavailable (select step failed)	Unavailable (select step failed)	2.700 \pm 0.200
50	Unavailable (select step failed)	Unavailable (select step failed)	12.480 \pm 0.208
100	Unavailable (select step failed)	Unavailable (select step failed)	99.787 \pm 2.113

1332 H.2 Comparison on Real Datasets.

1333 In this section, we compare RMSE and Max Variance on the real datasets: CPS, Adult, and Loans.
 1334 The different workloads are 1-way, 2-way, 3-way, 4-way, 5-way marginals, all ≤ 3 -way marginals,
 1335 and Small Marginals.

1336 H.2.1 RMSE Comparisons

1337 We provide an expanded comparison of RMSE on the 3 real datasets from the main paper. Here we
 1338 add more workloads. Table 16, 17 and 18 show the comparison of RMSE on the CPS, Adult, and
 1339 Loans datasets respectively.

1340 We notice that ResidualPlanner matches the theoretical SVD Bound while HDMM is slightly worse,
 1341 but still accurate. We conclude that when optimizing RMSE, the main advantage of ResidualPlanner
 1342 is superior scalability.

Table 16: Comparison of RMSE on CPS(5D) dataset.

Workload	HDMM	ResPlan	SVDB
1-way Marginals	1.756	1.744	1.744
2-way Marginals	2.103	2.035	2.035
3-way Marginals	2.089	2.048	2.048
4-way Marginals	1.648	1.627	1.627
5-way Marginals	1.000	1.000	1.000
≤ 3 -way Marginals	2.301	2.276	2.276
Small Marginals	2.525	2.525	2.525

Table 17: Comparison of RMSE on Adult(14D) dataset.

Workload	HDMM	ResPlan	SVDB
1-way Marginals	3.081	3.047	3.047
2-way Marginals	6.504	6.359	6.359
3-way Marginals	11.529	10.515	10.515
4-way Marginals	16.618	14.656	14.656
5-way Marginals	20.240	17.844	17.844
≤ 3 -way Marginals	11.555	10.665	10.665
Small Marginals	10.006	9.945	9.945

Table 18: Comparison of RMSE on Loans(12D) dataset.

Workload	HDMM	ResPlan	SVDB
1-way Marginals	2.903	2.875	2.875
2-way Marginals	5.747	5.634	5.634
3-way Marginals	9.478	8.702	8.702
4-way Marginals	12.537	11.267	11.267
5-way Marginals	14.872	12.678	12.678
≤ 3 -way Marginals	9.406	8.876	8.876
Small Marginals	8.262	8.206	8.206

1343 H.2.2 Max Variance

1344 The next comparison is on optimization for Max Variance. We repeat that HDMM only optimizes for
 1345 RMSE and this shows that optimizing for RMSE is highly suboptimal when one cares about max
 1346 variance.

1347 In contrast to RMSE, where the optimization problem generated by ResidualPlanner’s selection step
 1348 can be solved in closed form, for Max Variance, the optimization needs a convex solver. Hence we in-
 1349 clude comparisons between the open source ECOS [14] optimizer to the commercial Gurobi optimizer
 1350 [21]. Thus, our results have columns labeled ResidualPlanner+ECOS and ResidualPlanner+Gurobi.

1351 Tables 19, 20 and 21 show the results for the CPS, Adult, and Loans datasets, respectively. There is
 1352 one item to note about numerical stability. Although Gurobi is generally faster and more numerically
 1353 stable, the differences do not matter much. Situations where EOCS was worse are highlighted in red.
 1354 For example, in Table 19 for the CPS dataset, the dataset has only 5 attributes, so a 5-way marginal is
 1355 basically the entire dataset. The optimal mechanism for 5-way marginals simply adds $N(0, 1)$ noise
 1356 to each cell and optimizing for RMSE is equal to optimizing Max Variance for this special case. As
 1357 we see, the Max Variance for ResidualPlanner+ECOS is 1.008 which is 0.8% worse than optimal.
 1358 The reason for this is the numerical precision with which ECOS can solve the optimization problem
 1359 that ResidualPlanner gives it. In general, however, it looks like open source optimizers should work
 1360 fairly reliably for them to be used in real applications of ResidualPlanner.

Table 19: Comparison of Max Variance on CPS(5D) dataset.

Workload	HDMM	ResPlan + ECOS	ResPlan + Gurobi
1-way Marginals	13.672	4.346	4.346
2-way Marginals	47.741	7.897	7.897
3-way Marginals	71.549	7.706	7.706
4-way Marginals	15.538	4.142	4.141
5-way Marginals	1.000	1.008	1.000
≤ 3 -way Marginals	415.073	13.216	13.216
Small Marginals	223.579	11.774	11.774

Table 20: Comparison of Max Variance on Adult(14D) dataset.

Workload	HDMM	ResPlan + ECOS	ResPlan + Gurobi
1-way Marginals	41.772	12.047	12.047
2-way Marginals	599.843	67.802	67.802
3-way Marginals	5675.238	236.843	236.843
4-way Marginals	26959.322	575.213	575.213
5-way Marginals	79817.002	1030.948	1030.948
≤ 3 -way Marginals	6677.253	253.605	253.605
Small Marginals	2586.980	126.902	126.902

Table 21: Comparison of Max Variance on Loans(12D) dataset.

Workload	HDMM	ResPlan + ECOS	ResPlan + Gurobi
1-way Marginals	33.256	10.640	10.640
2-way Marginals	437.478	52.217	52.217
3-way Marginals	3095.997	156.638	156.638
4-way Marginals	13776.417	320.778	320.778
5-way Marginals	26056.289	474.244	474.243
\leq 3-way Marginals	4317.709	180.817	180.817
Small Marginals	2330.883	89.873	89.873