
Better Private Linear Regression Through Better Private Feature Selection

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Existing work on differentially private linear regression typically assumes that end
2 users can precisely set data bounds or algorithmic hyperparameters. End users
3 often struggle to meet these requirements without directly examining the data (and
4 violating privacy). Recent work has attempted to develop solutions that shift these
5 burdens from users to algorithms, but they struggle to provide utility as the feature
6 dimension grows. This work extends these algorithms to higher-dimensional
7 problems by introducing a differentially private feature selection method based on
8 Kendall rank correlation. We prove a utility guarantee for the setting where features
9 are normally distributed and conduct experiments across 25 datasets. We find that
10 adding this private feature selection step before regression significantly broadens
11 the applicability of “plug-and-play” private linear regression algorithms at little
12 additional cost to privacy, computation, or decision-making by the end user.

13 1 Introduction

14 Differentially private [10] algorithms employ carefully calibrated randomness to obscure the effect
15 of any single data point. Doing so typically requires an end user to provide bounds on input data to
16 ensure the correct scale of noise. However, end users often struggle to provide such bounds without
17 looking at the data itself [26], thus nullifying the intended privacy guarantee. This has motivated the
18 development of differentially private algorithms that do not require these choices from users.

19 To the best of our knowledge, two existing differentially private linear regression algorithms satisfy
20 this “plug-and-play” requirement: 1) the Tukey mechanism [2], which combines propose-test-release
21 with an exponential mechanism based on Tukey depth, and 2) Boosted AdaSSP [29], which applies
22 gradient boosting to the AdaSSP algorithm introduced by Wang [33]. We refer to these methods as,
23 respectively, Tukey and BAS. Neither algorithm requires data bounds, and both feature essentially
24 one chosen parameter (the number of models m for Tukey, and the number of boosting rounds T for
25 BAS) which admits simple heuristics without tuning. Tukey obtains strong empirical results when the
26 number of data points n greatly exceeds the feature dimension d [2], while BAS obtains somewhat
27 weaker performance on a larger class of datasets [29].

28 Nonetheless, neither algorithm provides generally strong utility on its own. Evaluated over a collection
29 of 25 linear regression datasets taken from Tang et al. [29], Tukey and BAS obtain coefficient of
30 determination $R^2 > 0$ on only four (see Section 4); for context, a baseline of $R^2 = 0$ is achieved by
31 the trivial constant predictor, which simply outputs the mean label. These results suggest room for
32 improvement for practical private linear regression.

33 1.1 Our Contributions

34 We extend existing work on private linear regression by adding a preprocessing step that applies
35 private feature selection. At a high level, this strategy circumvents the challenges of large feature
36 dimension d by restricting attention to $k \ll d$ carefully selected features. We initiate the study of
37 private feature selection in the context of “plug-and-play” private linear regression and make two
38 concrete contributions:

- 39 1. We introduce a practical algorithm, DPKendall, for differentially private feature selection
40 (Section 3). DPKendall uses Kendall rank correlation [19] and only requires the user to
41 choose the number k of features to select. It satisfies ϵ -DP and, given n samples with
42 d -dimensional features, runs in time $O(dkn \log(n))$ (Theorem 3.4). We also provide a
43 utility guarantee when the features are normally distributed (Theorem 3.8).
- 44 2. We conduct experiments across 25 datasets (Section 4), with k fixed at 5 and 10. These
45 compare Tukey and BAS without feature selection, with SubLasso feature selection [20],
46 and with DPKendall feature selection. Using $(\ln(3), 10^{-5})$ -DP to cover both private feature
47 selection and private regression, we find at $k = 5$ that adding DPKendall yields $R^2 > 0$ on
48 56% of the datasets. Replacing DPKendall with SubLasso drops the rate to 40% of datasets,
49 and omitting feature selection entirely drops it further to 16%.

50 In summary, we suggest that DPKendall significantly expands the applicability and utility of practical
51 private linear regression.

52 1.2 Related Work

53 The focus of this work is practical private feature selection applied to private linear regression. We
54 therefore refer readers interested in a more general overview of the private linear regression literature
55 to the discussions of Amin et al. [2] and Tang et al. [29].

56 Several works have studied private sparse linear regression [20, 30, 16, 28]. However, Jain and
57 Thakurta [16] and Talwar et al. [28] require an ℓ_∞ bound on the input data, and the stability test that
58 powers the feature selection algorithm of Thakurta and Smith [30] requires the end user to provide
59 granular details about the optimal Lasso model. These requirements are significant practical obstacles.
60 An exception is the work of Kifer et al. [20]. Their algorithm first performs feature selection using
61 subsampling and aggregation of non-private Lasso models. This feature selection method, which
62 we call SubLasso, only requires the end user to select the number of features k . To the selected
63 features, Kifer et al. [20] then apply objective perturbation to privately optimize the Lasso objective.
64 As objective perturbation requires the end user to choose parameter ranges and provides a somewhat
65 brittle privacy guarantee contingent on the convergence of the optimization, we do not consider it
66 here. Instead, our experiments combine SubLasso feature selection with the Tukey and BAS private
67 regression algorithms. An expanded description of SubLasso appears in Section 4.2.

68 We now turn to the general problem of private feature selection. There is a significant literature
69 studying private analogues of the general technique of principal component analysis (PCA) [24, 14,
70 8, 18, 11, 1]. Unfortunately, all of these algorithms assume some variant of a bound on the row norm
71 of the input data. Stoddard et al. [27] studied private feature selection in the setting where features
72 and labels are binary, but it is not clear how to extend their methods to the non-binary setting that we
73 consider in this work. SubLasso is therefore the primary comparison private feature selection method
74 in this paper. We are not aware of existing work that studies private feature selection in the specific
75 context of “plug-and-play” private linear regression.

76 Finally, private rank correlation has previously been studied by Kusner et al. [22]. They derived a
77 different, normalized sensitivity bound appropriate for their “swap” privacy setting and applied it to
78 privately determine the causal relationship between two random variables.

79 2 Preliminaries

80 Throughout this paper, a database D is a collection of labelled points (x, y) where $x \in \mathbb{R}^d$, $y \in \mathbb{R}$,
81 and each user contributes a single point. We use the “add-remove” form of differential privacy.

82 **Definition 2.1** ([10]). Databases D, D' from data domain \mathcal{D} are neighbors $D \sim D'$ if they differ
83 in the presence or absence of a single record. A randomized mechanism $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{O}$ is (ε, δ) -
84 differentially private (DP) if for all $D \sim D' \in \mathcal{D}$ and any $S \subseteq \mathcal{O}$

$$\mathbb{P}_{\mathcal{M}}[\mathcal{M}(D) \in S] \leq e^\varepsilon \mathbb{P}_{\mathcal{M}}[\mathcal{M}(D') \in S] + \delta.$$

85 When $\delta = 0$, we say \mathcal{M} is ε -DP.

86 We use basic composition to reason about the privacy guarantee obtained from repeated application of
87 a private algorithm. More sophisticated notions of composition exist, but for our setting of relatively
88 few compositions, basic composition is simpler and suffers negligible utility loss.

89 **Lemma 2.2** ([10]). Suppose that for $j \in [k]$, algorithm \mathcal{A}_j is $(\varepsilon_j, \delta_j)$ -DP. Then running all k
90 algorithms is $(\sum_{j=1}^k \varepsilon_j, \sum_{j=1}^k \delta_j)$ -DP.

91 Both SubLasso and DPKendall use a private subroutine for identifying the highest count item(s) from
92 a collection, known as private top- k . Several algorithms for this problem exist [4, 9, 13]. We use the
93 pure DP “peeling mechanism” based on Gumbel noise [9], as its analysis is relatively simple, and its
94 performance is essentially identical to other variants for the relatively small k used in this paper.

95 **Definition 2.3** ([9]). A Gumbel distribution with parameter b is defined over $x \in \mathbb{R}$ by $\mathbb{P}[x; b] =$
96 $\frac{1}{b} \cdot \exp(-\frac{x}{b} - e^{-x/b})$. Given $c = (c_1, \dots, c_d) \in \mathbb{R}^d$, $k \in \mathbb{N}$, and privacy parameter ε ,
97 $\text{Peel}(c, k, \Delta_\infty, \varepsilon)$ adds independent Gumbel $(\frac{2k\Delta_\infty}{\varepsilon})$ noise to each count c_j and outputs the or-
98 dered sequence of indices with the largest noisy counts.

99 **Lemma 2.4** ([9]). Given $c = (c_1, \dots, c_d) \in \mathbb{R}^d$ with ℓ_∞ sensitivity Δ_∞ , $\text{Peel}(c, k, \Delta_\infty, \varepsilon)$ is ε -DP.

100 The primary advantage of Peel over generic noise addition is that, although users may contribute to d
101 counts, the added noise only scales with k . We note that while Peel requires an ℓ_∞ bound, neither
102 DPKendall nor SubLasso needs user input to set it: regardless of the dataset, DPKendall’s use of
103 Peel has $\ell_\infty = 3$ and SubLasso’s use has $\ell_\infty = 1$ (see Algorithms 1 and 2).

104 3 Feature Selection Algorithm

105 This section describes our feature selection algorithm, DPKendall, and formally analyzes its utility.
106 Section 3.1 introduces and discusses Kendall rank correlation, Section 3.2 describes the full algorithm,
107 and the utility result appears in Section 3.3.

108 3.1 Kendall Rank Correlation

109 The core statistic behind our algorithm is Kendall rank correlation. Informally, Kendall rank correla-
110 tion measures the strength of a monotonic relationship between two variables.

111 **Definition 3.1** ([19]). Given a collection of data points $(X, Y) = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$ and
112 $i < i'$, a pair of observations $(X_i, Y_i), (X_{i'}, Y_{i'})$ is discordant if $(X_i - X_{i'})(Y_i - Y_{i'}) < 0$. Given
113 data (X, Y) , let $d_{X,Y}$ denote the number of discordant pairs. Then the empirical Kendall rank
114 correlation is

$$\hat{\tau}(X, Y) := \frac{n}{2} - \frac{2d_{X,Y}}{n-1}.$$

115 For real random variables X and Y , we can also define the population Kendall rank correlation by

$$\tau(X, Y) = \mathbb{P}[(X - X')(Y - Y') > 0] - \mathbb{P}[(X - X')(Y - Y') < 0].$$

116 Kendall rank correlation is therefore high when an increase in X or Y typically accompanies an
117 increase in the other, low when an increase in one typically accompanies a decrease in the other,
118 and close to 0 when a change in one implies little about the other. We typically focus on empirical
119 Kendall rank correlation, but the population definition will be useful in the proof of our utility result.

120 Before discussing Kendall rank correlation in the context of privacy, we note two straightforward
121 properties. First, for simplicity, we use a version of Kendall rank correlation that does not account for
122 ties. We ensure this in practice by perturbing data by a small amount of continuous random noise.
123 Second, τ has range $[-1, 1]$, but (this paper’s version of) $\hat{\tau}$ has range $[-n/2, n/2]$. This scaling does
124 not affect the qualitative interpretation and ensures that $\hat{\tau}$ has low sensitivity¹.

¹In particular, without scaling, $\hat{\tau}$ would be $\approx \frac{1}{n}$ -sensitive, but n is private information in add-remove privacy.

125 **Lemma 3.2.** $\Delta(\hat{\tau}) = 3/2$.

126 *Proof.* At a high level, the proof verifies that the addition or removal of a user changes the first term
127 of $\hat{\tau}$ by at most $1/2$, and the second term by at most 1.

128 In more detail, consider two neighboring databases (X, Y) and (X', Y') . Without loss
129 of generality, we may assume that $(X, Y) = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$ and $(X', Y') =$
130 $\{(X'_1, Y'_1), \dots, (X'_{n+1}, Y'_{n+1})\}$ where for all $i \in [n]$ we have that $X'_i = X_i$ and $Y'_i = Y_i$. First, we
131 argue that the number of discordant pairs in (X', Y') cannot be much larger than in (X, Y) . By
132 definition, we have that $d_{X', Y'} - d_{X, Y} = \sum_{j=1}^n \mathbb{1}_{(X_j - X'_{n+1})(Y_j - Y'_{n+1}) < 0}$. In particular, this implies
133 that $d_{X', Y'} - d_{X, Y} \in [0, n]$.

134 We can rewrite the difference in Kendall correlation between (X, Y) and (X', Y') as follows:

$$\begin{aligned} \hat{\tau}(X, Y) - \hat{\tau}(X', Y') &= \frac{n}{2} - \frac{2d_{X, Y}}{n-1} - \frac{n+1}{2} + \frac{2d_{X', Y'}}{n} \\ &= 2 \left(\frac{d_{X', Y'}}{n} - \frac{d_{X, Y}}{n-1} \right) - \frac{1}{2} \\ &= 2 \left(\frac{d_{X', Y'} - d_{X, Y}}{n} + \frac{d_{X, Y}}{n} - \frac{d_{X, Y}}{n-1} \right) - \frac{1}{2} \\ &= 2 \frac{d_{X', Y'} - d_{X, Y}}{n} + 2d_{X, Y} \left(\frac{1}{n} - \frac{1}{n-1} \right) - \frac{1}{2} \\ &= 2 \frac{d_{X', Y'} - d_{X, Y}}{n} - \frac{d_{X, Y}}{\binom{n}{2}} - \frac{1}{2}, \end{aligned}$$

135 where the final equality follows from the fact that $2(\frac{1}{n} - \frac{1}{n-1}) = -1/\binom{n}{2}$. Using our previous
136 calculation, the first term is in the range $[0, 2]$ and, since $d_{X, Y} \in [0, \binom{n}{2}]$, the second term is in the
137 range $[-1, 0]$. It follows that

$$-\frac{3}{2} = 0 - 1 - \frac{1}{2} \leq \hat{\tau}(X, Y) - \hat{\tau}(X', Y') \leq 2 - 0 - \frac{1}{2} = \frac{3}{2},$$

138 and therefore $|\hat{\tau}(X, Y) - \hat{\tau}(X', Y')| \leq 3/2$.

139 To show that the sensitivity is not smaller than $3/2$, consider neighboring databases (X, Y) and
140 (X', Y') such that $d_{X, Y} = 0$ and (X', Y') contains a new point that is discordant with all points in
141 (X, Y) . Then $d_{X', Y'} = n$ while $d_{X, Y} = 0$. Then $\hat{\tau}(X, Y) - \hat{\tau}(X', Y') = 2 - 0 - 1/2 = 3/2$. \square

142 Turning to privacy, Kendall rank correlation has two notable strengths. First, because it is computed
143 entirely from information about the relative ordering of data, it does not require an end user to provide
144 data bounds. This makes it a natural complement to private regression methods that also operate
145 without user-provided data bounds. Second, Kendall rank correlation's sensitivity is constant, but
146 its range scales linearly with n . This makes it easy to compute privately. A contrasting example is
147 Pearson correlation, which requires data bounds to compute covariances and has sensitivity identical
148 to its range. An extended discussion of alternative notions of correlation appears in Section 6.

149 Finally, Kendall rank correlation can be computed relatively quickly using a variant of merge sort.

150 **Lemma 3.3 ([21]).** *Given collection of data points $(X, Y) = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$, $\hat{\tau}(X, Y)$
151 can be computed in time $O(n \log(n))$.*

152 3.2 DPKendall

153 Having defined Kendall rank correlation, we now describe our private feature selection algorithm,
154 DPKendall. Informally, DPKendall balances two desiderata: 1) selecting features that correlate with
155 the label, and 2) selecting features that do not correlate with previously selected features. Prioritizing
156 only the former selects for redundant copies of a single informative feature, while prioritizing only
157 the latter selects for features that are pure noise.

158 In more detail, DPKendall consists of k applications of Peel to select a feature that is correlated with
159 the label and relatively uncorrelated with the features already chosen. Thus, letting S_t denote the set

160 of $t - 1$ features already chosen in round t , each round attempts to compute

$$\max_{j \notin S_t} \left(|\hat{\tau}(X_j, Y)| - \frac{1}{t-1} \sum_{j' \in S_t} |\hat{\tau}(X_j, X_{j'})| \right). \quad (1)$$

161 The $\frac{1}{t-1}$ scaling ensures that the sensitivity of the overall quantity remains fixed at $\frac{3}{2}$ in the first round
 162 and 3 in the remaining rounds. Note that in the first round we take second term to be 0, and only label
 163 correlation is considered.

Algorithm 1 DPKendall(D, k, ε)

1: **Input:** Examples $D = \{(X_i, Y_i)\}_{i=1}^n$, number of selected features k , privacy parameter ε
 2: **for** $j = 1, \dots, d$ **do**
 3: Compute $\hat{\tau}_j^Y = |\hat{\tau}(X_j, Y)|$
 4: Initialize $S = \emptyset$
 5: Initialize $\hat{\tau} = \hat{\tau}^Y \in \mathbb{R}^d$
 6: Initialize $\hat{\tau}^X = 0 \in \mathbb{R}^d$
 7: **for** $t = 1, \dots, k$ **do**
 8: Set $\Delta_\infty = \frac{3}{2} + \frac{3}{2} \cdot \mathbb{1}_{t>1}$
 9: Set $s_t = \text{Peel} \left(\hat{\tau}^Y - \frac{\hat{\tau}^X}{t-1}, 1, \Delta_\infty, \frac{\varepsilon}{k} \right)$
 10: Expand $S = S \cup s_t$
 11: Update $\hat{\tau}_{s_t}^Y = -\infty$
 12: **for** $j \notin S$ **do**
 13: Update $\hat{\tau}_j^X = \hat{\tau}_j^X - |\hat{\tau}(X_j, X_{s_t})|$
 14: **Return** S

164 Pseudocode for DPKendall appears in Algorithm 1. Its runtime and privacy are easy to verify.

165 **Theorem 3.4.** DPKendall runs in time $O(dkn \log(n))$ and satisfies ε -DP.

166 *Proof.* By Lemma 3.3, each computation of Kendall rank correlation takes time $O(n \log(n))$, so
 167 Line 2’s loop takes time $O(dn \log(n))$, as does each execution of Line 12’s loop. Each call to Peel
 168 requires $O(d)$ samples of Gumbel noise and thus contributes $O(dk)$ time overall. The loop in Line 7
 169 therefore takes time $O(dkn \log(n))$. The privacy guarantee follows from Lemmas 2.4 and 3.2. \square

170 For comparison, standard OLS on n samples of data with d features requires time $O(d^2n)$; DPKendall
 171 is asymptotically no slower as long as $n \leq O(2^{d/k})$. Since we typically take $k \ll d$, DPKendall is
 172 computationally “free” in many realistic data settings.

173 3.3 Utility Guarantee

174 The proof of DPKendall’s utility guarantee combines results about population Kendall rank correla-
 175 tion (Lemma 3.5), empirical Kendall rank correlation concentration (Lemma 3.6), and the accuracy
 176 of Peel (Lemma 3.7). The final guarantee (Theorem 3.8) demonstrates that DPKendall selects useful
 177 features even in the presence of redundant features.

178 We start with the population Kendall rank correlation guarantee. Its proof, and all proofs for uncited
 179 results in this section, appears in Section 7 in the Appendix.

180 **Lemma 3.5.** Suppose that X_1, \dots, X_k are independent random variables where $X_j \sim N(\mu_j, \sigma_j^2)$.
 181 Let $\xi \sim N(0, \sigma_e^2)$ be independent noise. Then if the label is generated by $Y = \sum_{j=1}^k \beta_j X_j + \xi$, for
 182 any $j^* \in [k]$,

$$\tau(X_{j^*}, Y) = \frac{2}{\pi} \cdot \arctan \frac{\beta_{j^*} \sigma_{j^*}}{\sqrt{\sum_{j \neq j^*} \beta_j^2 \sigma_j^2 + \sigma_e^2}}.$$

183 To interpret this result, recall that $\tau \in [-1, 1]$ and arctan has domain \mathbb{R} , is odd, and has
 184 $\lim_{x \rightarrow \infty} \arctan x = \pi/2$. Lemma 3.5 thus says that if we fix the other σ_j and σ_e and take $\sigma_{j^*} \rightarrow \infty$,
 185 $\tau(X_{j^*}, Y) \rightarrow \text{sign}(\beta_{j^*})$ as expected. The next step is to verify that $\hat{\tau}$ concentrates around τ .

186 **Lemma 3.6** (Lemma 1 [3]). Given n observations each of random variables X and Y , with proba-
 187 bility $1 - \eta$,

$$|\hat{\tau}(X, Y) - \frac{n}{2} \cdot \tau(X, Y)| \leq \sqrt{8n \ln(2/\eta)}.$$

188 Finally, we state a basic accuracy result for the Gumbel noise employed by Peel.

189 **Lemma 3.7.** Given i.i.d. random variables $X_1, \dots, X_d \sim \text{Gumbel}(b)$, with probability $1 - \eta$,

$$\max_{j \in [d]} |X_j| \leq b \ln \left(\frac{2d}{\eta} \right).$$

190 We now have the tools necessary for the final result.

191 **Theorem 3.8.** Suppose that X_1, \dots, X_k are independent random variables where each $X_j \sim$
 192 $N(\mu_j, \sigma_j^2)$. Suppose additionally that of the remaining $d - k$ random variables, for each $j \in [k]$,
 193 n_j are copies of X_j , where $\sum_{j=1}^k n_j \leq d - k$. For each $j \in [k]$, let S_j denote the set of indices
 194 consisting of j and the indices of its copies. Then if the label is generated by $Y = \sum_{j=1}^k \beta_j X_j + \xi$
 195 where $\xi \sim N(0, \sigma_e^2)$ is independent random noise, if

$$n = \Omega \left(\frac{k \cdot \ln(dk/\eta)}{\varepsilon \cdot \min_{j^* \in [k]} \left\{ \left| \arctan \frac{\beta_{j^*} \sigma_{j^*}}{\sqrt{\sum_{j \neq j^*} \beta_j^2 \sigma_j^2 + \sigma_e^2}} \right| \right\}} \right),$$

196 then with probability $1 - O(\eta)$, DPKendall correctly selects exactly one index from each of S_1, \dots, S_k .

197 *Proof.* The proof reduces to applying the preceding lemmas with appropriate union bounds. Dropping
 198 the constant scaling of η for neatness, with probability $1 - O(\eta)$:

199 **1.** Feature-label correlations are large for informative features and small for uninformative features:
 200 by Lemma 3.5 and Lemma 3.6, each feature in $j^* \in \cup_{j \in [k]} S_j$ has

$$\hat{\tau}(X_{j^*}, Y) \geq \frac{n}{\pi} \cdot \arctan \frac{\beta_{j^*} \sigma_{j^*}}{\sqrt{\sum_{j \neq j^*} \beta_j^2 \sigma_j^2 + \sigma_e^2}} - \sqrt{8n \ln(d/\eta)}$$

201 and any $j^* \notin \cup_{j \in [k]} S_j$ has $\hat{\tau}(X_{j^*}, Y) \leq \sqrt{8n \ln(d/\eta)}$.

202 **2.** Feature-feature correlations are large between copies of a feature and small between independent
 203 features: by Lemma 3.5 and Lemma 3.6, for any $j \in [k]$ and $j_1, j_2 \in S_j$,

$$\hat{\tau}(X_{j_1}, X_{j_2}) \geq \frac{n}{2} - \sqrt{8n \ln(d/\eta)}$$

204 and for any j_1, j_2 such that there exists no S_j containing both, $\hat{\tau}(X_{j_1}, X_{j_2}) \leq \sqrt{8n \ln(d/\eta)}$.

205 **3.** The at most dk draws of Gumbel noise have absolute value bounded by $\frac{k}{\varepsilon} \ln \left(\frac{dk}{\eta} \right)$.

206 Combining these results, to ensure that DPKendall's k calls to Peel produce exactly one index from
 207 each of S_1, \dots, S_k , it suffices to have

$$n \cdot \min_{j^* \in [k]} \left\{ \left| \arctan \frac{\beta_{j^*} \sigma_{j^*}}{\sqrt{\sum_{j \neq j^*} \beta_j^2 \sigma_j^2 + \sigma_e^2}} \right| \right\} = \Omega \left(\left[\sqrt{n} + \frac{k}{\varepsilon} \right] \ln \left(\frac{dk}{\eta} \right) \right)$$

208 which rearranges to yield the claim. \square

209 4 Experiments

210 This section collects experimental evaluations of DPKendall and other methods on 25 of the 33
 211 datasets² used by Tang et al. [29]. Descriptions of the relevant algorithms appear in Section 4.1 and
 212 Section 4.2. Section 4.3 discusses the results.

²We omit the datasets with OpenML task IDs 361080-361084 as they are restricted versions of other included datasets. We also exclude 361090 and 361097 as non-private OLS obtained $R^2 \ll 0$ on both. Details for the remaining datasets appear in Figure 3 in the Appendix.

213 **4.1 Feature Selection Baseline**

214 Our experiments use SubLasso [20] as a baseline “plug-and-play” private feature selection method.
 215 At a high level, the algorithm randomly partitions its data into m subsets, computes a non-private
 216 Lasso regression model on each, and then privately aggregates these models to select k significant
 217 features. The private aggregation process is simple; for each subset’s learned model, choose the k
 218 features with largest absolute coefficient, then apply private top- k to compute the k features most
 219 selected by the m models. Kifer et al. [20] introduced and analyzed SubLasso; we collect its relevant
 220 properties in Lemma 4.1. Pseudocode appears in Algorithm 2.

Algorithm 2 SubLasso(D, k, m, ε)

- 1: **Input:** Examples $D = \{(X_i, Y_i)\}_{i=1}^n$, number of selected features k , number of models m ,
 privacy parameter ε
 - 2: Randomly partition D into m equal-size subsets S_1, \dots, S_m
 - 3: **for** $i = 1, \dots, m$ **do**
 - 4: Compute Lasso model θ_i on S_i
 - 5: Compute set C_i of the k indices of θ_i with largest absolute value
 - 6: Compute binary vector $v_i \in \{0, 1\}^d$ where $v_{i,j} = \mathbb{1}_{j \in C_i}$
 - 7: Compute $V \in \mathbb{R}^d = \sum_{i=1}^m v_i$
 - 8: Return Peel($V, k, 1, \varepsilon$)
-

221 **Lemma 4.1.** SubLasso is ε -DP and runs in time $O(d^2n)$.

222 *Proof.* The privacy guarantee is immediate from that of Peel (Lemma 2.4). Solving Lasso on n/m
 223 data points with d -dimensional features takes time $O(\frac{d^2n}{m})$ [12]. Multiplying through by m produces
 224 the final result, since Peel only takes time $O(dk)$. \square

225 Finally, we briefly discuss the role of the intercept feature in SubLasso. As with all algorithms in
 226 our experiments, we add an intercept feature (with a constant value of 1) to each vector of features.
 227 Each Lasso model is trained on data with this intercept. However, the intercept is removed before the
 228 private voting step, k features are chosen from the remaining features, and the intercept is added back
 229 afterward. This ensures that privacy is not wasted on the intercept feature, which we always include.

230 **4.2 Comparison Algorithms**

231 We evaluate seven algorithms:

- 232 1. NonDP is a non-private baseline running generic ordinary least-squares regression.
- 233 2. BAS runs Boosted AdaSSP [29] without feature selection. We imitate the parameter settings
 234 used by Tang et al. [29] and set feature and gradient clipping norms to 1 and the number of
 235 boosting rounds to 100 throughout.
- 236 3. Tukey runs the Tukey mechanism [2] without feature selection. We introduce and use
 237 a tighter version of the propose-test-release (PTR) check given by Amin et al. [2]. This
 238 reduces the number of models needed for PTR to pass. The proof appears in Section 9 in the
 239 Appendix and may be of independent interest. To privately choose the number of models m
 240 used by the Tukey mechanism, we first privately estimate a $1 - \eta$ probability lower bound on
 241 the number of points n using the Laplace CDF, $\tilde{n} = n + \text{Lap}(\frac{1}{\varepsilon'}) - \frac{\ln(1/2\eta)}{\varepsilon'}$, and then set
 242 the number of models to $m = \lfloor \tilde{n}/d \rfloor$. Tukey spends 5% of its ε privacy budget estimating
 243 m and the remainder on the Tukey mechanism.
- 244 4. L-BAS runs spends 5% of its ε privacy budget choosing $m = \lfloor \tilde{n}/k \rfloor$ for SubLasso, 5% of ε
 245 running SubLasso, and then spends the remainder to run BAS on the selected features.
- 246 5. L-Tukey spends 5% of its ε privacy budget choosing $m = \lfloor \tilde{n}/k \rfloor$ for SubLasso, 5% running
 247 SubLasso, and the remainder running Tukey on the selected features using the same m .
- 248 6. K-BAS spends 5% of its ε privacy budget running DPKendall and then spends the remainder
 249 running BAS on the selected features.
- 250 7. K-Tukey spends 5% of its ε privacy budget choosing $m = \lfloor \tilde{n}/k \rfloor$, 5% running DPKendall,
 251 and then spends the remainder running the Tukey mechanism on the selected features.

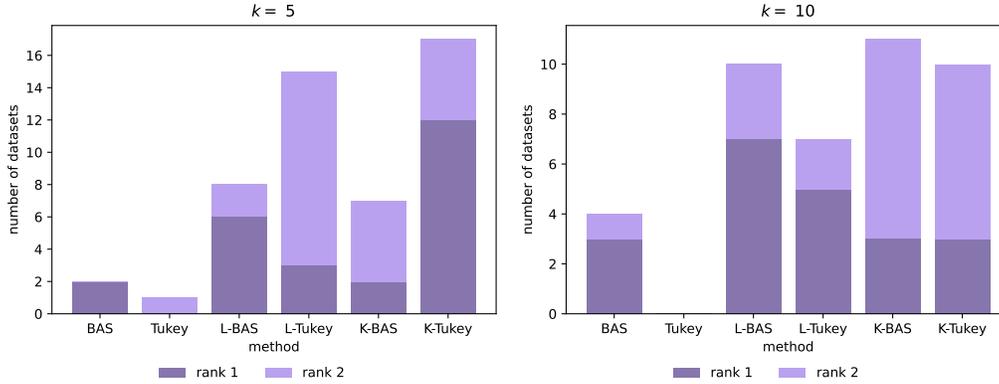


Figure 1: Plots of rank data for each private method.

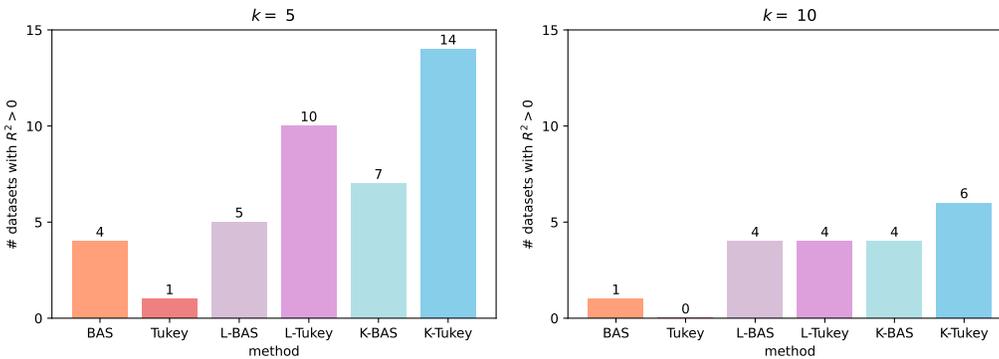


Figure 2: Plots of the number of datasets with positive R^2 for each private method.

252 4.3 Results

253 All experiments use $(\ln(3), 10^{-5})$ -DP. Where applicable, 5% of the privacy budget is spent on private
 254 feature selection, 5% on choosing the number of models, and the remainder is spent on private
 255 regression. Throughout, we use $\eta = 10^{-4}$ as the failure probability for the lower bound used to
 256 choose the number of models. For each algorithm and dataset, we run 10 trials using random 90-10
 257 train-test splits and record the resulting test R^2 values. Tables of the results at $k = 5$ and $k = 10$
 258 appear in Section 10 in the Appendix. A condensed presentation appears below. At a high level, we
 259 summarize the results in terms of *relative* and *absolute* performance.

260 4.3.1 Relative Performance

261 First, for each dataset and method, we compute the median R^2 of the final model across the 10 trials
 262 and then rank the methods, with the best method receiving a rank of 1. Figure 1 plots the number of
 263 times each method is ranked first or second.

264 At $k = 5$ (left), K-Tukey performs best by a significant margin: it ranks first on 48% of datasets,
 265 twice the fraction of any other method. It also ranks first or second on the largest fraction of datasets
 266 (68%). In some contrast, L-BAS obtains the top ranking on 24% of datasets, whereas K-BAS only
 267 does so on 8%; nonetheless, the two have nearly the same number of total first or second rankings.
 268 At $k = 10$ (right), no clear winner emerges among the feature selecting methods, as L-BAS, K-BAS,
 269 and K-Tukey are all first or second on around half the datasets³, though the methods using SubLasso
 270 have a higher share of datasets ranked first.

³Note that $k = 10$ only uses 21 datasets. This is because 4 of the 25 datasets used for $k = 5$ have $d < 10$.

271 **4.3.2 Absolute Performance**

272 Second, we count the number of datasets on which the method achieves a positive median R^2
273 (Figure 2), recalling that $R^2 = 0$ is achieved by the trivial model that always predicts the mean
274 label. The $k = 5$ (left) setting again demonstrates clear trends: K-Tukey attains $R^2 > 0$ on 56%
275 of datasets, L-Tukey does so on 40%, K-BAS does so on 28%, and L-BAS on 20%. DPKendall
276 therefore consistently demonstrates stronger performance than SubLasso. As in the rank data, at
277 $k = 10$ the picture is less clear. However, K-Tukey is still best by some margin, with the remaining
278 feature selecting methods all performing roughly equally well.

279 **4.4 Discussion**

280 A few trends are apparent from Section 4.3. First, DPKendall generally achieves stronger final utility
281 than SubLasso, particularly for the Tukey mechanism; the effect is similar but smaller for $k = 10$;
282 and feature selection generally improves the performance of private linear regression.

283 **Comparing Feature Selection Algorithms.** A possible reason for DPKendall’s improvement over
284 SubLasso is that, while SubLasso takes advantage of the stability properties that Lasso exhibits in
285 certain data regimes [20], this stability does not always hold in practice. Another possible explanation
286 is that the feature coefficients passed to Peel scale with $O(n)$ for DPKendall and $m = O(\frac{n}{d})$ or
287 $m = O(\frac{n}{k})$ for SubLasso. Both algorithm’s invocations of Peel add noise scaling with $O(\frac{k}{\epsilon})$, so
288 DPKendall’s larger scale makes it more robust to privacy-preserving noise. Finally, we emphasize
289 that DPKendall achieves this even though its $O(dkn \log(n))$ runtime is asymptotically smaller than
290 the $O(d^2n)$ runtime of SubLasso in most settings.

291 **Choosing k .** Next, we examine the decrease in performance from $k = 5$ to $k = 10$. Conceptually,
292 past a certain point adding marginally less informative features to a private model may worsen utility
293 due to the privacy cost of considering these features. Moving from $k = 5$ to $k = 10$ may cross this
294 threshold for many of our datasets; note from Figure 4 and Figure 5 in the Appendix that, of the 21
295 datasets used for $k = 10$, 86% witness their highest private R^2 in the $k = 5$ setting⁴. Moreover, from
296 $k = 5$ to $k = 10$ the total number of positive R^2 datasets across methods declines by more than 50%,
297 from 41 to 19, with all method achieving positive R^2 less frequently at $k = 10$ than $k = 5$. We
298 therefore suggest $k = 5$ as the more relevant setting, and a good choice in practice.

299 **The Effect of Private Feature Selection.** Much work aims to circumvent generic lower bounds for
300 privately answering queries by taking advantage of instance-specific structure [6, 15, 32, 5]. Similar
301 works exist for private optimization, either by explicitly incorporating problem information [34, 17]
302 or showing that problem-agnostic algorithms can, under certain conditions, take advantage of problem
303 structure organically [23]. We suggest that this paper makes a similar contribution: feature selection
304 reduces the need for algorithms like Boosted AdaSSP and the Tukey mechanism to “waste” privacy
305 on computations over irrelevant features. This enables them to apply less obscuring noise to the
306 signal contained in the selected features. The result is the significant increase in utility shown here.

307 **5 Conclusion**

308 We briefly discuss a few of DPKendall’s limitations. First, it requires an end user to choose the
309 number of features k to select. Second, DPKendall’s use of Kendall rank correlation may struggle
310 when ties are intrinsic to the data’s structure, e.g., when the data is categorical, as a monotonic
311 relationship between feature and label becomes less applicable. Finally, Kendall rank correlation may
312 fail to distinguish between a feature with a strong linear monotonic relationship with the label and
313 a feature with a strong nonlinear monotonic relationship with the label, even though the former is
314 more likely to be useful for linear regression. Unfortunately, it is not obvious how to incorporate
315 relationships more sophisticated than simple monotonicity without sacrificing rank correlation’s low
316 sensitivity. Answering these questions may be an interesting avenue for future work.

317 Nonetheless, the results of this paper demonstrate that DPKendall expands the applicability of plug-
318 and-play private linear regression algorithms while providing more utility in less time than the current
319 state of the art. We therefore suggest that DPKendall presents a step forward for practical private
320 linear regression.

⁴The exceptions are datasets 361075, 361091, and 361103.

References

- 321
- 322 [1] Kareem Amin, Travis Dick, Alex Kulesza, Andres Munoz Medina, and Sergei Vassilvitskii.
323 [Differentially private covariance estimation](#). *Neural Information Processing Systems (NeurIPS)*,
324 2019.
- 325 [2] Kareem Amin, Matthew Joseph, Mónica Ribero, and Sergei Vassilvitskii. [Easy differentially
326 private linear regression](#). In *International Conference on Learning Representations (ICLR)*,
327 2023.
- 328 [3] Youssef Anzarmou, Abdallah Mkhadri, and Karim Oualkacha. [The Kendall interaction filter for
329 variable interaction screening in high dimensional classification problems](#). *Journal of Applied
330 Statistics*, 2022.
- 331 [4] Raghav Bhaskar, Srivatsan Laxman, Adam Smith, and Abhradeep Thakurta. [Discovering
332 frequent patterns in sensitive data](#). In *Knowledge Discovery and Data Mining (KDD)*, 2010.
- 333 [5] Jaroslaw Błasiok, Mark Bun, Aleksandar Nikolov, and Thomas Steinke. [Towards instance-
334 optimal private query release](#). In *Symposium on Discrete Algorithms (SODA)*, 2019.
- 335 [6] Avrim Blum, Katrina Ligett, and Aaron Roth. [A learning theory approach to non-interactive
336 database privacy](#). In *Symposium on the Theory of Computing (STOC)*, 2008.
- 337 [7] Gavin Brown, Marco Gaboardi, Adam Smith, Jonathan Ullman, and Lydia Zakyntinou.
338 [Covariance-Aware Private Mean Estimation Without Private Covariance Estimation](#). In *Neural
339 Information Processing Systems (NeurIPS)*, 2021.
- 340 [8] Kamalika Chaudhuri, Anand Sarwate, and Kaushik Sinha. [Near-optimal differentially private
341 principal components](#). *Neural Information Processing Systems (NIPS)*, 2012.
- 342 [9] David Durfee and Ryan M Rogers. [Practical differentially private top-k selection with pay-
343 what-you-get composition](#). *Neural Information Processing Systems (NeurIPS)*, 2019.
- 344 [10] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. [Calibrating noise to
345 sensitivity in private data analysis](#). In *Theory of Cryptography Conference (TCC)*, 2006.
- 346 [11] Cynthia Dwork, Kunal Talwar, Abhradeep Thakurta, and Li Zhang. [Analyze gauss: optimal
347 bounds for privacy-preserving principal component analysis](#). In *Symposium on Theory of
348 Computing (STOC)*, 2014.
- 349 [12] Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. [Least angle regression](#).
350 *Annals of Statistics*, 2004.
- 351 [13] Jennifer Gillenwater, Matthew Joseph, Andres Munoz Medina, and Monica Ribero Diaz. [A
352 Joint Exponential Mechanism For Differentially Private Top-k](#). In *International Conference on
353 Machine Learning (ICML)*, 2022.
- 354 [14] Moritz Hardt and Aaron Roth. [Beating randomized response on incoherent matrices](#). In
355 *Symposium on Theory of Computing (ICML)*, 2012.
- 356 [15] Moritz Hardt and Guy N Rothblum. [A multiplicative weights mechanism for privacy-preserving
357 data analysis](#). In *Foundations of Computer Science (FOCS)*, 2010.
- 358 [16] Prateek Jain and Abhradeep Guha Thakurta. [\(Near\) dimension independent risk bounds for
359 differentially private learning](#). In *International Conference on Machine Learning (ICML)*, 2014.
- 360 [17] Peter Kairouz, Monica Ribero Diaz, Keith Rush, and Abhradeep Thakurta. [\(Nearly\) Dimension
361 Independent Private ERM with AdaGrad Rates via Publicly Estimated Subspaces](#). In *Conference
362 on Learning Theory (COLT)*, 2021.
- 363 [18] Michael Kapralov and Kunal Talwar. [On differentially private low rank approximation](#). In
364 *Symposium on Discrete Algorithms (SODA)*, 2013.
- 365 [19] Maurice G Kendall. [A new measure of rank correlation](#). *Biometrika*, 1938.

- 366 [20] Daniel Kifer, Adam Smith, and Abhradeep Thakurta. [Private convex empirical risk minimization](#)
367 [and high-dimensional regression](#). In *Conference on Learning Theory (COLT)*, 2012.
- 368 [21] William R Knight. [A computer method for calculating Kendall’s tau with ungrouped data](#).
369 *Journal of the American Statistical Association*, 1966.
- 370 [22] Matt J Kusner, Yu Sun, Karthik Sridharan, and Kilian Q Weinberger. [Private causal inference](#).
371 In *Artificial Intelligence and Statistics (AISTATS)*, 2016.
- 372 [23] Xuechen Li, Daogao Liu, Tatsunori B Hashimoto, Huseyin A Inan, Janardhan Kulkarni, Yin-Tat
373 Lee, and Abhradeep Guha Thakurta. [When Does Differentially Private Learning Not Suffer in](#)
374 [High Dimensions?](#) *Neural Information Processing Systems (NeurIPS)*, 2022.
- 375 [24] Frank McSherry and Ilya Mironov. [Differentially private recommender systems: Building](#)
376 [privacy into the netflix prize contenders](#). In *Knowledge Discovery and Data Mining (KDD)*,
377 2009.
- 378 [25] Donald Bruce Owen. [A table of normal integrals](#). *Communications in Statistics-Simulation and*
379 *Computation*, 1980.
- 380 [26] Jayshree Sarathy, Sophia Song, Audrey Emma Haque, Tania Schlatter, and Salil Vadhan. [Don’t](#)
381 [look at the data! How differential privacy reconfigures data subjects, data analysts, and the](#)
382 [practices of data science](#). In *Conference on Human Factors in Computing Systems (CHI)*, 2023.
- 383 [27] Ben Stoddard, Yan Chen, and Ashwin Machanavajjhala. [Differentially private algorithms for](#)
384 [empirical machine learning](#). *arXiv preprint arXiv:1411.5428*, 2014.
- 385 [28] Kunal Talwar, Abhradeep Guha Thakurta, and Li Zhang. [Nearly optimal private lasso](#). *Neural*
386 *Information Processing Systems (NIPS)*, 2015.
- 387 [29] Shuai Tang, Sergul Aydore, Michael Kearns, Saeyoung Rho, Aaron Roth, Yichen Wang, Yu-
388 Xiang Wang, and Zhiwei Steven Wu. [Improved Differentially Private Regression via Gradient](#)
389 [Boosting](#). *arXiv preprint arXiv:2303.03451*, 2023.
- 390 [30] Abhradeep Guha Thakurta and Adam Smith. [Differentially private feature selection via stability](#)
391 [arguments, and the robustness of the lasso](#). In *Conference on Learning Theory (COLT)*, 2013.
- 392 [31] John W Tukey. Mathematics and the picturing of data. In *International Congress of Mathemati-*
393 *cians (IMC)*, 1975.
- 394 [32] Jonathan Ullman. [Private multiplicative weights beyond linear queries](#). In *Principles of*
395 *Database Systems (PODS)*, 2015.
- 396 [33] Yu-Xiang Wang. [Revisiting differentially private linear regression: optimal and adaptive](#)
397 [prediction & estimation in unbounded domain](#). In *Uncertainty in Artificial Intelligence (UAI)*,
398 2018.
- 399 [34] Yingxue Zhou, Zhiwei Steven Wu, and Arindam Banerjee. [Bypassing the ambient dimension:](#)
400 [Private sgd with gradient subspace identification](#). *International Conference on Learning*
401 *Representations (ICLR)*, 2021.

402 **6 Alternative Notions of Correlation**

403 **6.1 Pearson**

404 The Pearson correlation between some feature X and the label Y is defined by

$$r(X, Y) := \frac{\text{Cov}(X, Y)}{\sqrt{\text{Var}(X)\text{Var}(Y)}}. \quad (2)$$

405 Evaluated on a sample of n points, this becomes

$$r(X, Y) := \frac{\sum_{i=1}^n (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^n (X_i - \bar{X})^2} \sqrt{\sum_{i=1}^n (Y_i - \bar{Y})^2}}.$$

406 where \bar{X} and \bar{Y} are sample means.

407 **Lemma 6.1.** $r \in [-1, 1]$.

408 *Proof.* This is immediate from Cauchy-Schwarz. □

409 Note that a value of -1 is perfect anticorrelation and a value of 1 is perfect correlation. A downside
410 of Pearson correlation is that it is not robust. In particular, its sensitivity is the same as its range.

411 **Lemma 6.2.** *Pearson correlation has ℓ_1 -sensitivity $\Delta(r) = 2$.*

412 *Proof.* Consider neighboring databases $D = \{(-1, -1), (1, 1)\}$ and $D' =$
413 $\{(-1, -1), (1, 1), (-c, c)\}$. $r(D) = 1$, but

$$\begin{aligned} r(D') &= \frac{(-1 + \frac{c}{3})(-1 - \frac{c}{3}) + (1 + \frac{c}{3})(1 - \frac{c}{3}) + (-\frac{2c}{3} \cdot \frac{2c}{3})}{\sqrt{(-1 + \frac{c}{3})^2 + (1 + \frac{c}{3})^2 + (-\frac{2c}{3})^2} \sqrt{(-1 - \frac{c}{3})^2 + (1 - \frac{c}{3})^2 + (\frac{2c}{3})^2}} \\ &\approx \frac{-(\frac{2c}{3})^2}{(\frac{2c}{3})^2} = -1 \end{aligned}$$

414 where the approximation is increasingly accurate as c grows. □

415 **6.2 Spearman**

416 Spearman rank correlation is Pearson correlation applied to rank variables.

417 **Definition 6.3.** *Given data points $X_1, \dots, X_n \in \mathbb{R}$, the corresponding rank variables*
418 *$R(X_1), \dots, R(X_n)$ are defined by setting $R(X_i)$ to the position of X_i when X_1, \dots, X_n are sorted in*
419 *descending order. Given data (X, Y) , the Spearman rank correlation is $\rho(X, Y) := r(R(X), R(Y))$.*

420 For example, given database $D = \{(0, 1), (1, 0), (2, 3)\}$, its Spearman rank correlation is

$$\rho(D) := r(\{(3, 2), (2, 3), (1, 1)\}).$$

421 A useful privacy property of rank is that it does not depend on data scale. Moreover, if there are no
422 ties then Spearman rank correlation admits a simple closed form.

423 **Lemma 6.4.** $\rho(X, Y) = 1 - \frac{6 \sum_{i=1}^n (R(X_i) - R(Y_i))^2}{n(n^2 - 1)}$.

424 If we consider adding a “perfectly unsorted” data point with rank variables $(1, n + 1)$ to a perfectly
425 sorted database with rank variables $\{(1, 1), (2, 2), \dots, (n, n)\}$, ρ changes from 1 to $1 - \frac{6(n+n^2)}{n(n^2-1)}$.

426 The sensitivity’s dependence on n complicates its usage with add-remove privacy. Nonetheless, both
427 Spearman and Kendall correlation’s use of rank makes them relatively easy to compute privately, and
428 as the two methods are often used interchangeably in practice, we opt for Kendall rank correlation for
429 simplicity.

430 7 Deferred Proofs From Section 3.3

431 We first restate and prove Lemma 3.5.

432 **Lemma 3.5.** *Suppose that X_1, \dots, X_k are independent random variables where $X_j \sim N(\mu_j, \sigma_j^2)$.
433 Let $\xi \sim N(0, \sigma_e^2)$ be independent noise. Then if the label is generated by $Y = \sum_{j=1}^k \beta_j X_j + \xi$, for
434 any $j^* \in [k]$,*

$$\tau(X_{j^*}, Y) = \frac{2}{\pi} \cdot \arctan \frac{\beta_{j^*} \sigma_{j^*}}{\sqrt{\sum_{j \neq j^*} \beta_j^2 \sigma_j^2 + \sigma_e^2}}.$$

435 *Proof.* Recall from Definition 3.1 the population formulation of τ ,

$$\tau(X, Y) = \mathbb{P}[(X - X')(Y - Y') > 0] - \mathbb{P}[(X - X')(Y - Y') < 0] \quad (3)$$

436 where X and X' are i.i.d., as are Y and Y' . In our case, we can define $Z = \sum_{j \neq j^*} \beta_j X_j + \xi$ and
437 rewrite the first term of Equation 3 for our setting as

$$\int_0^\infty f_{X_{j^*} - X'_{j^*}}(t) \cdot [1 - F_{Z - Z'}(-\beta_{j^*} t)] dt + \int_{-\infty}^0 f_{X_{j^*} - X'_{j^*}}(t) \cdot F_{Z - Z'}(-\beta_{j^*} t) dt. \quad (4)$$

438 where f_A and F_A denote densities and cumulative distribution functions of random variable A ,
439 respectively. Since the relevant distributions are all Gaussian, $X_{j^*} - X'_{j^*} \sim N(0, 2\sigma_{j^*}^2)$ and $Z - Z' \sim$
440 $N(0, 2[\sum_{j \neq j^*} \beta_j^2 \sigma_j^2 + \sigma_e^2])$. For neatness, shorthand $\sigma^2 = 2\sigma_{j^*}^2$ and $\sigma_{-1}^2 = 2[\sum_{j \neq j^*} \beta_j^2 \sigma_j^2 + \sigma_e^2]$.
441 Then if we let ϕ denote the PDF of a standard Gaussian and Φ the CDF of a standard Gaussian,
442 Equation 4 becomes

$$\begin{aligned} & \int_0^\infty \frac{1}{\sigma} \phi(t/\sigma) \cdot [1 - \Phi(-\beta_{j^*} t/\sigma_{-1})] dt + \int_{-\infty}^0 \frac{1}{\sigma} \phi(t/\sigma) \Phi(-\beta_{j^*} t/\sigma_{-1}) dt \\ &= \frac{1}{\sigma} \left[\int_0^\infty \phi(t/\sigma) \Phi(\beta_{j^*} t/\sigma_{-1}) dt + \int_{-\infty}^0 \phi(-t/\sigma) \Phi(-\beta_{j^*} t/\sigma_{-1}) dt \right] \\ &= \frac{2}{\sigma} \int_0^\infty \phi(t/\sigma) \Phi(\beta_{j^*} t/\sigma_{-1}) dt. \end{aligned}$$

443 We can similarly analyze the second term of Equation 3 to get

$$\begin{aligned} & \int_0^\infty f_{X_{j^*} - X'_{j^*}}(t) \cdot F_{Z - Z'}(-\beta_{j^*} t) dt + \int_{-\infty}^0 f_{X_{j^*} - X'_{j^*}}(t) \cdot [1 - F_{Z - Z'}(-\beta_{j^*} t)] dt \\ &= \int_0^\infty \frac{1}{\sigma} \phi(t/\sigma) \Phi(-\beta_{j^*} t/\sigma_{-1}) dt + \int_{-\infty}^0 \frac{1}{\sigma} \phi(t/\sigma) \Phi(\beta_{j^*} t/\sigma_{-1}) dt \\ &= \frac{2}{\sigma} \int_{-\infty}^0 \phi(t/\sigma) \Phi(\beta_{j^*} t/\sigma_{-1}) dt. \end{aligned}$$

444 Using both results, we get

$$\begin{aligned} \tau(X_1, Y) &= \frac{2}{\sigma} \left[\int_0^\infty \phi(t/\sigma) \Phi(\beta_{j^*} t/\sigma_{-1}) dt - \int_{-\infty}^0 \phi(t/\sigma) \Phi(\beta_{j^*} t/\sigma_{-1}) dt \right] \\ &= \frac{2}{\sigma} \left[\int_0^\infty \phi(t/\sigma) \Phi(\beta_{j^*} t/\sigma_{-1}) dt - \int_0^\infty \phi(t/\sigma) [1 - \Phi(\beta_{j^*} t/\sigma_{-1})] dt \right] \\ &= \frac{4}{\sigma} \int_0^\infty \phi(t/\sigma) \Phi(\beta_{j^*} t/\sigma_{-1}) dt - 1 \\ &= \frac{4}{\sigma} \cdot \frac{\sigma}{2\pi} \left(\frac{\pi}{2} + \arctan \frac{\beta_{j^*} \sigma}{\sigma_{-1}} \right) - 1 \\ &= \frac{2}{\pi} \cdot \arctan \frac{\beta_{j^*} \sigma}{\sigma_{-1}}. \end{aligned}$$

445 where the third equality uses $\int_0^\infty \phi(t/\sigma) dt = \frac{\sigma}{2}$ and the fourth equality comes from Equation 1,010.4
446 of Owen [25]. Substituting in the values of σ and σ_{-1} yields the claim. \square

447 Next is Lemma 3.7.

448 **Lemma 3.7.** *Given i.i.d. random variables $X_1, \dots, X_d \sim \text{Gumbel}(b)$, with probability $1 - \eta$,*

$$\max_{j \in [d]} |X_j| \leq b \ln \left(\frac{2d}{\eta} \right).$$

449 *Proof.* Recall from Definition 2.3 that Gumbel(b) has density $f(x) = \frac{1}{b} \cdot \exp(-\frac{x}{b} - e^{-x/b})$. Then
 450 $\frac{f(x)}{f(-x)} = \exp(-\frac{x}{b} - e^{-x/b} - \frac{x}{b} + e^{x/b})$. For $z \geq 0$, by $e^z = \sum_{n=0}^{\infty} \frac{z^n}{n!}$,

$$2z + e^{-z} \leq 2z + 1 - z + \frac{z^2}{2} = 1 + z + \frac{z^2}{2} \leq e^z$$

451 so since $b \geq 0$, $f(x) \geq f(-x)$ for $x \geq 0$. Letting $F(z) = \exp(-\exp(-z/b))$ denote the CDF for
 452 Gumbel(b), it follows that for $z \geq 0$, $1 - F(z) \geq F(-z)$. Thus the probability that $\max_{j \in [d]} |X_j|$
 453 exceeds t is upper bounded by $2d(1 - F(t))$. The claim then follows from rearranging the inequality

$$\begin{aligned} 2d(1 - F(t)) &\leq \eta \\ 1 - \frac{\eta}{2d} &\leq F(t) \\ 1 - \frac{\eta}{2d} &\leq \exp(-\exp(-t/b)) \\ \ln \left(1 - \frac{\eta}{2d} \right) &\leq -\exp(-t/b) \\ -b \ln \left(-\ln \left(1 - \frac{\eta}{2d} \right) \right) &\leq t \\ b \ln \left(\frac{1}{-\ln(1 - \eta/[2d])} \right) &\leq t \end{aligned}$$

454 and using $-\ln(1 - x) = \sum_{i=1}^{\infty} x^i/i \geq x$. □

455 8 Datasets

456 A summary of the datasets used in our experiments appears in Figure 3.

457 We briefly discuss the role of the intercept in these datasets. Throughout, we explicitly add an
 458 intercept feature (constant 1) to each vector. Where feature selection is applied, we explicitly remove
 459 the intercept feature during feature selection and then add it back to the k selected features afterward.
 460 The resulting regression problem therefore has dimension $k + 1$. We do this to avoid spending privacy
 461 budget selecting the intercept feature.

462 9 Modified PTR Lemma

463 This section describes a simple tightening of Lemma 3.6 from Amin et al. [2] (which is itself a
 464 small modification of Lemma 3.8 from Brown et al. [7]). Tukey uses the result as its propose-test-
 465 release (PTR) check, so tightening it makes the check easier to pass. Proving the result will require
 466 introducing details of the Tukey algorithm. The following exposition aims to keep this document both
 467 self-contained and brief; the interested reader should consult the expanded treatment given by Amin
 468 et al. [2] for further details.

469 Tukey depth was introduced by Tukey [31]. Amin et al. [2] used an approximation for efficiency.
 470 Roughly, Tukey depth is a notion of depth for a collection of points in space. (Exact) Tukey depth is
 471 evaluated over all possible directions in \mathbb{R}^d , while approximate Tukey depth is evaluated only over
 472 axis-aligned directions.

473 **Definition 9.1** ([31, 2]). *A halfspace h_v is defined by a vector $v \in \mathbb{R}^d$, $h_v = \{y \in \mathbb{R}^d \mid \langle v, y \rangle \geq 0\}$.
 474 Let $E = \{e_1, \dots, e_d\}$ be the canonical basis for \mathbb{R}^d and let $D \subset \mathbb{R}^d$. The approximate Tukey depth of*

OpenML Task ID	n	d	$\lfloor n/d \rfloor$	$\lfloor n/5 \rfloor$	$\lfloor n/10 \rfloor$
361072	8192	22	372	1638	819
361073	15000	27	555	3000	1500
361074	16599	17	976	3319	1659
361075	7797	614	12	1559	779
361076	6497	12	541	1299	649
361077	13750	34	404	2750	1375
361078	20640	9	2293	4128	2064
361079	22784	17	1340	4556	2278
361085	10081	7	1440	2016	1008
361087	13932	14	995	2786	1393
361088	21263	80	265	4252	2126
361089	20640	9	2293	4128	2064
361091	515345	91	5663	103069	51534
361092	8885	83	107	1777	888
361093	4052	13	311	810	405
361094	8641	6	1440	1728	864
361095	166821	24	6950	33364	16682
361096	53940	27	1997	10788	5394
361098	10692	18	594	2138	1069
361099	17379	21	827	3475	1737
361100	39644	74	535	7928	3964
361101	581835	32	18182	116367	58183
361102	21613	20	1080	4322	2161
361103	394299	27	14603	78859	39429
361104	241600	16	15100	48320	24160

Figure 3: Parameters of the 25 datasets used in our experiments.

475 a point $y \in \mathbb{R}^d$ with respect to D , denoted $\tilde{T}_D(y)$, is the minimum number of points in D in any of
476 the 2d halfspaces determined by E containing y ,

$$\tilde{T}_D(y) = \min_{v \in \pm E \text{ s.t. } y \in h_v} \sum_{x \in D} \mathbb{1}_{x \in h_v}.$$

477 At a high level, Lemma 3.6 from Amin et al. [2] is a statement about the volumes of regions of
478 different depths. The next step is to formally define these volumes.

479 **Definition 9.2** ([2]). Given database D , define $S_{i,D} = \{y \in \mathbb{R}^d \mid \tilde{T}_D(y) \geq i\}$ to be the set of
480 points with approximate Tukey depth at least i in D and $V_{i,D} = \text{vol}(S_{i,D})$ to be the volume of
481 that set. When D is clear from context, we write S_i and V_i for brevity. We also use $w_D(V_{d,D}) :=$
482 $\int_{S_{d,D}} \exp(\varepsilon \cdot \tilde{T}_D(y)) dy$ to denote the weight assigned to $V_{d,D}$ by an exponential mechanism whose
483 score function is \tilde{T}_D .

484 Amin et al. [2] define a family of mechanisms A_1, A_2, \dots , where A_t runs the exponential mechanism
485 to choose a point of approximately maximal Tukey depth, but restricted to the domain of points with
486 Tukey depth at least t . Since this domain is a data-dependent quantity, they use the PTR framework
487 to select a safe depth t . We briefly recall the definitions of “safe” and “unsafe” databases given
488 by Brown et al. [7], together with the key PTR result from Amin et al. [2].

489 **Definition 9.3** (Definitions 2.1 and 3.1 [7]). Two distributions \mathcal{P}, \mathcal{Q} over domain \mathcal{W} are (ε, δ) -
490 indistinguishable, denoted $\mathcal{P} \approx_{\varepsilon, \delta} \mathcal{Q}$, if for any measurable subset $W \subset \mathcal{W}$,

$$\mathbb{P}_{w \sim \mathcal{P}} [w \in W] \leq e^\varepsilon \mathbb{P}_{w \sim \mathcal{Q}} [w \in W] + \delta \text{ and } \mathbb{P}_{w \sim \mathcal{Q}} [w \in W] \leq e^\varepsilon \mathbb{P}_{w \sim \mathcal{P}} [w \in W] + \delta.$$

491 Database D is (ε, δ, t) -safe if for all neighboring $D' \sim D$, we have $A_t(D) \approx_{\varepsilon, \delta} A_t(D')$. Let
492 $\text{Safe}_{(\varepsilon, \delta, t)}$ be the set of safe databases, and let $\text{Unsafe}_{(\varepsilon, \delta, t)}$ be its complement.

493 We can now restate Lemma 3.6 from Amin et al. [2]. Informally, it states that if the volume of an
494 “outer” region of Tukey depth is not much larger than the volume of an “inner” region, the difference

495 in depth between the two is a lower bound on the distance to an unsafe database. For the purpose of
 496 this paper, Tukey applies this result by finding such a k , adding noise for privacy, and checking that
 497 the resulting distance to unsafety is large enough that subsequent steps will be privacy-safe.

498 **Lemma 9.4.** Define $M(D)$ to be a mechanism that receives as input database D and computes the
 499 largest $k \in \{0, \dots, t-1\}$ such that there exists $g > 0$ where

$$\frac{V_{t-k-1,D}}{V_{t+k+g+1,D}} \cdot e^{-\varepsilon g/2} \leq \delta$$

500 or outputs -1 if the inequality does not hold for any such k . Then for arbitrary D

- 501 1. M is 1-sensitive, and
- 502 2. for all $z \in \text{Unsafe}_{(\varepsilon, 4e^\varepsilon \delta, t)}$, $d_H(D, z) > M(D)$.

503 We provide a drop-in replacement for Lemma 9.4 that slightly weakens the requirement placed on k .

504 **Lemma 9.5.** Define $M(D)$ to be a mechanism that receives as input database D and computes the
 505 largest $k \in \{0, \dots, t-1\}$ such that

$$\frac{V_{t-k-1,D}}{w_D(V_{t+k-1,D})} \cdot e^{\varepsilon(t+k+1)} \leq \delta$$

506 or outputs -1 if the inequality does not hold for any such k . Then for arbitrary D

- 507 1. M is 1-sensitive, and
- 508 2. for all $z \in \text{Unsafe}_{(\varepsilon, 4e^\varepsilon \delta, t)}$, $d_H(D, z) > M(D)$.

509 The new result therefore replaces the denominator $V_{t+k+g+1,D} \cdot e^{\varepsilon g/2}$ with denominator $\frac{w_D(V_{t+k-1,D})}{e^{\varepsilon(t+k+1)}}$.
 510 Every point in $V_{t+k-1,D}$ of depth at least $t+k+g+1$ has score at least $t+k+g+1$, so
 511 $\frac{w_D(V_{t+k-1,D})}{e^{\varepsilon(t+k+1)}} \geq V_{t+k+g+1,D} \cdot e^{\varepsilon g}$, so $V_{t+k+g+1,D} \cdot e^{\varepsilon g/2} \leq \frac{w_D(V_{t+k-1,D})}{e^{\varepsilon(t+k+1)}}$ and the check for the
 512 new result is no harder to pass. To see that it may be easier, note that only the new result takes
 513 advantage of the higher scores of deeper points in $V_{t+k-1,D}$.

514 *Proof of Lemma 9.5.* First we prove item 1. Let D and D' be any neighboring databases and suppose
 515 WLOG that $D' = D \cup \{x\}$. We want to show that $|M(D) - M(D')| \leq 1$.

516 First we prove relationships between the points with approximate Tukey depth at least p and $p-1$ in
 517 datasets D and D' . From the definition of approximate Tukey depth, together with the fact that D'
 518 contains one additional point, for any point y , we are guaranteed that $\tilde{T}_D(y) \leq \tilde{T}_{D'}(y) \leq \tilde{T}_D(y) + 1$.
 519 Recall that $S_{p,D}$ is the set of points with approximate Tukey depth at least p in D . This implies that
 520 $S_{p+1,D} \subset S_{p,D}$. Next, since for every point y we have $\tilde{T}_{D'}(y) \geq \tilde{T}_D(y)$, we have that $S_{p,D} \subset S_{p,D'}$.
 521 Finally, since $\tilde{T}_D(y) \geq \tilde{T}_{D'}(y) - 1$, we have that $S_{p,D'} \subset S_{p-1,D}$. Taken together, we have

$$S_{p+1,D} \subset S_{p,D} \subset S_{p,D'} \subset S_{p-1,D}.$$

522 It follows that

$$V_{p+1,D} \leq V_{p,D} \leq V_{p,D'} \leq V_{p-1,D}. \quad (5)$$

523 Next, since the unnormalized exponential mechanism density $y \mapsto \exp(y\tilde{T}_D(y))$ is non-negative,
 524 we have that $w_D(V_{p+1,D}) = \int_{S_{p+1,D}} \exp(\varepsilon\tilde{T}_D(y)) dy \leq \int_{S_{p,D}} \exp(\varepsilon\tilde{T}_D(y)) dy = w_D(V_{p,D})$.
 525 Using the fact that $\tilde{T}_D(y) \leq \tilde{T}_{D'}(y)$, we have that $w_D(V_{p,D}) = \int_{S_{p,D}} \exp(\varepsilon\tilde{T}_D(y)) dy \leq$
 526 $\int_{S_{p,D'}} \exp(\varepsilon\tilde{T}_{D'}(y)) dy = w_{D'}(V_{p,D'})$. Finally, using the fact that $\tilde{T}_{D'}(y) \leq \tilde{T}_D(y) + 1$, we have
 527 $w_{D'}(V_{p,D'}) = \int_{S_{p,D'}} \exp(\varepsilon\tilde{T}_{D'}(y)) dy \leq \int_{S_{p-1,D}} \exp(\varepsilon\tilde{T}_D(y) + \varepsilon) dy = e^\varepsilon w_D(V_{p-1,D})$. Together,
 528 this gives

$$w_D(V_{p+1,D}) \leq w_D(V_{p,D}) \leq w_{D'}(V_{p,D'}) \leq w_D(V_{p-1,D}) \cdot e^\varepsilon. \quad (6)$$

529 Now suppose there exists $k_{D'}^* \geq 0$ such that $\frac{V_{t-k_{D'}^*-1,D'}}{w_{D'}(V_{t+k_{D'}^*+2,D'})} \cdot e^{\varepsilon(t+k_{D'}^*+1)} \leq \delta$. Then by Equa-
 530 tion 5, $V_{t-k_{D'}^*-1,D'} \geq V_{t-k_{D'}^*,D}$, and by Equation 6 $w_{D'}(V_{t+k_{D'}^*+2,D'}) \leq w_D(V_{t+k_{D'}^*+1,D}) \cdot e^\varepsilon$, so

531 $\frac{V_{t-k_{D'}^*,D}}{w_D(V_{t+k_{D'}^*,D})} \cdot e^{\varepsilon(t+k_{D'}^*)} \leq \delta$ and then $k_D^* \geq k_{D'}^* - 1$. Similarly, if there exists $k_D^* \geq 0$ such that
532 $\frac{V_{t-k_D^*-1,D}}{w_D(V_{t+k_D^*+2,D})} \cdot e^{\varepsilon(t+k_D^*+1)} \leq \delta$, then by Equation 5 $V_{t-k_D^*-1,D} \geq V_{t-k_D^*,D'}$, and by Equation 6
533 $w_D(V_{t+k_D^*+2,D}) \leq w_{D'}(V_{t+k_D^*+2,D'})$, so $\frac{V_{t-k_D^*,D'}}{w_{D'}(V_{t+k_D^*+2,D'})} \cdot e^{\varepsilon(t+k_{D'}^*)} \leq \delta$, and $k_{D'}^* \geq k_D^* - 1$. Thus
534 if $k_D^* \geq 0$ or $k_{D'}^* \geq 0$, $|k_D^* - k_{D'}^*| \leq 1$. The result then follows since $k^* \geq -1$.

535 As in Lemma 9.4, item 2 is a consequence of Lemma 3.8 from Brown et al. [7] and the fact that
536 $k^* = -1$ is a trivial lower bound on distance. The only change made to the proof of Lemma 3.8
537 of Brown et al. [7] is, in its notation, to replace its denominator lower bound with

$$w_z(V_{t-1,z}) \geq w_x(V_{t+k-1,x}) \cdot e^{-k\varepsilon}$$

538 This uses the fact that x and z differ in at most the addition or removal of k data points. Thus
539 $V_{t+k-1,x} \leq V_{t-1,z}$, and no point's score increases by more than k from $V_{t-1,z}$ to $V_{t+k-1,x}$. Since
540 their numerator upper bound is $V_{t-k-1,x} \cdot e^{\varepsilon(t+1)}$ (note that the 2 is dropped here because approximate
541 Tukey depth is monotonic; see Section 7.3 of [2] for details), the result follows. \square

542 10 Extended Experiment Results

Task ID	NonDP	BAS	Tukey	L-BAS	L-Tukey	K-BAS	K-Tukey
361072	7.3e-01	7.8e-01	$-\infty$	-2.7e-02	1.0e-01	2.1e-01	1.0e-01
361073	4.6e-01	-3.4e+00	$-\infty$	-4.6e-01	-2.3e-02	-4.8e-01	-4.3e-01
361074	8.0e-01	-4.5e+04	-6.6e+02	-8.2e+02	-1.7e+02	-5.6e+05	3.1e-01
361075	6.2e-01	-3.6e+02	$-\infty$	3.2e-03	2.2e-02	5.2e-03	7.1e-02
361076	2.8e-01	-3.4e+00	$-\infty$	-5.6e-01	-2.5e-01	-4.3e+00	8.5e-02
361077	8.2e-01	-2.1e+08	$-\infty$	-1.0e+07	3.7e-01	-1.9e+08	6.6e-01
361078	6.5e-01	-1.1e+03	-6.4e+00	-7.4e-01	-1.2e+00	-6.1e+02	-9.5e-01
361079	2.6e-01	-5.4e+00	-1.0e+01	-1.3e+01	-2.e+00	-9.4e+00	-6.8e-01
361085	3.3e-01	-2.e+01	3.3e-01	-1.7e+01	2.8e-01	-1.3e+01	3.7e-01
361087	7.2e-01	-3.1e+03	-8.8e+04	-4.7e+00	-1.9e+02	-2.8e+00	6.6e-01
361088	7.3e-01	5.4e-02	$-\infty$	1.4e-01	3.e-01	2.1e-01	3.6e-01
361089	6.2e-01	-3.e+00	-1.5e+01	-1.3e+00	-5.9e+00	-5.9e+04	-2.6e-01
361091	2.4e-01	-3.0e+04	-1.9e+00	-3.0e+04	3.8e-02	-3.1e+04	4.9e-02
361092	2.0e-02	-4.7e+05	$-\infty$	-4.1e+04	-1.4e+08	-1.3e+05	-9.3e+08
361093	4.3e-01	-5.6e-02	$-\infty$	-4.6e-02	$-\infty$	-4.9e-02	$-\infty$
361094	8.3e-01	4.2e-01	-3.8e+05	-1.0e+00	-4.2e+05	-2.4e+00	-3.0e+05
361095	2.2e-01	-1.0e+00	-1.2e+09	-8.0e-01	2.e-01	1.5e-01	2.1e-01
361096	9.7e-01	-9.1e+02	-1.e+08	-4.9e+00	9.1e-01	7.1e-01	8.8e-01
361098	8.6e-01	-2.9e+01	$-\infty$	-3.9e+00	-3.1e+00	-2.2e-01	-5.6e+00
361099	4.0e-01	-4.4e-01	-4.8e+10	-4.8e-01	-9.8e-03	-4.2e-01	3.2e-01
361100	1.2e-01	-6.3e+01	$-\infty$	-1.1e+00	-2.8e-02	-4.6e-01	-1.7e-01
361101	3.3e-01	-9.3e+01	-5.9e+08	3.6e-01	2.e-01	3.0e-01	2.1e-01
361102	7.6e-01	-4.1e+02	-1.4e+15	6.8e-02	-2.6e-01	-4.0e+00	-5.3e+00
361103	5.3e-01	4.3e-01	-6.3e+07	5.e-01	4.9e-01	3.3e-01	4.8e-01
361104	6.8e-01	-3.7e+03	-6.4e+07	-1.5e+02	-2.6e+05	-8.3e-01	-2.9e+01

Figure 4: R^2 values for $k = 5$. Each entry is the median test value from 10 trials. The top two private values for each dataset are bolded.

Task ID	NonDP	BAS	Tukey	L-BAS	L-Tukey	K-BAS	K-Tukey
361072	7.4e-01	5.3e-01	$-\infty$	3.5e-01	1.8e-01	7.3e-01	8.2e-02
361073	4.6e-01	-8.2e-01	$-\infty$	-4.4e-01	-1.7e+00	-4.8e-01	-1.6e+01
361074	8.1e-01	-2.3e+06	-1.3e+03	-5.6e+06	-3.4e+02	-7.2e+04	-1.4e+02
361075	6.2e-01	-1.7e+02	$-\infty$	7.5e-02	-9.7e-01	3.1e-02	-2.8e+00
361076	2.8e-01	-2.4e+02	-8.2e+04	-2.7e+01	-3.5e+03	-2.4e+01	-1.4e+03
361077	7.9e-01	-1.1e+08	$-\infty$	-1.4e+07	-3.1e+01	-2.1e+08	-1.7e+02
361079	2.4e-01	-1.8e+00	-3.7e+01	-1.2e+01	-8.1e-01	-4.1e+00	-1.5e+00
361087	7.1e-01	-3.7e+01	-1.5e+04	-7.2e+00	-4.8e+03	-8.5e+01	-2.e+01
361088	7.3e-01	-1.9e+00	$-\infty$	1.2e-01	9.2e-02	1.2e-01	7.1e-02
361091	2.4e-01	-3.0e+04	-3.4e+00	-3.0e+04	6.8e-02	-3.1e+04	5.7e-02
361092	4.0e-02	-2.4e+06	$-\infty$	-8.7e+05	-7.6e+09	-7.5e+05	-3.7e+10
361093	4.3e-01	-3.2e-02	$-\infty$	-5.e-01	$-\infty$	-1.1e-01	$-\infty$
361095	2.2e-01	-1.9e+01	-1.2e+09	-8.6e-01	-2.4e+07	-1.4e+00	2.1e-01
361096	9.8e-01	-3.7e+02	-1.0e+08	-3.4e+03	-1.8e+00	-1.0e+03	4.2e-01
361098	8.6e-01	-1.8e+02	$-\infty$	-4.1e-01	-1.3e+07	-9.5e-01	-5.8e-01
361099	4.0e-01	-4.3e-01	-2.3e+10	-4.9e-01	-2.4e+09	-4.8e-01	-2.5e+08
361100	1.2e-01	-9.2e+01	$-\infty$	-1.6e+00	-8.2e-02	-3.5e+00	-5.2e-01
361101	3.4e-01	-3.3e+03	-1.7e+08	-8.1e-01	-1.4e+02	-9.5e-01	-4.7e+05
361102	7.5e-01	-9.5e+02	-1.7e+15	-1.3e+00	-1.5e+15	-2.5e+01	-1.4e+05
361103	5.4e-01	-2.6e+00	-1.2e+08	4.9e-01	5.1e-01	3.5e-01	4.9e-01
361104	6.8e-01	-9.e+02	-1.0e+08	-2.9e+03	-3.5e+07	-1.4e+03	-6.6e+05

Figure 5: This figure records the same information as Figure 4, but for $k = 10$.