



Epistemic Parity: Reproducibility as an Evaluation Metric for Differential Privacy

Lucas Rosenblatt*
New York University
New York, NY, USA
lucas.rosenblatt@nyu.edu

Bernease Herman
University of Washington
Seattle, WA, USA
bernease@uw.edu

Anastasia Holovenko
Ukrainian Catholic
University
Lviv, Ukraine
anastasia.holovenko@ucu.edu.ua

Wonkwon Lee
New York University
New York, NY, USA
wl2733@nyu.edu

Joshua Loftus
London School of
Economics
London, UK
J.R.Loftus@lse.ac.uk

Elizabeth McKinnie
Microsoft
Seattle, WA, USA
Elizabeth.McKinnie@microsoft.com

Taras Rumezhak
Ukrainian Catholic
University
Lviv, Ukraine
rumezhak@ucu.edu.ua

Andrii Stadnik
Ukrainian Catholic
University
Lviv, Ukraine
andrii.stadnik@ucu.edu.ua

Bill Howe
University of Washington
Seattle, WA, USA
billhowe@uw.edu

Julia Stoyanovich
New York University
New York, NY, USA
stoyanovich@nyu.edu

ABSTRACT

Differential privacy (DP) data synthesizers are increasingly proposed to afford public release of sensitive information, offering theoretical guarantees for privacy (and, in some cases, utility), but limited empirical evidence of utility in practical settings. Utility is typically measured as the error on representative proxy tasks, such as descriptive statistics, multivariate correlations, the accuracy of trained classifiers, or performance over a query workload. The ability for these results to generalize to practitioners' experience has been questioned in a number of settings, including the U.S. Census. In this paper, we propose an evaluation methodology for synthetic data that avoids assumptions about the representativeness of proxy tasks, instead measuring the likelihood that published conclusions would change had the authors used synthetic data, a condition we call epistemic parity. Our methodology consists of reproducing empirical conclusions of peer-reviewed papers on real, publicly available data, then re-running these experiments a second time on DP synthetic data and comparing the results.

We instantiate our methodology over a benchmark of recent peer-reviewed papers that analyze public datasets in the ICPSR social science repository. We model quantitative claims computationally to automate the experimental workflow, and model qualitative claims by reproducing visualizations and comparing the results manually. We then generate DP synthetic datasets using multiple state-of-the-art mechanisms, and estimate the likelihood that these conclusions will hold. We find that, for reasonable privacy regimes, state-of-the-art DP synthesizers are able to achieve high epistemic parity for several papers in our benchmark. However, some papers, and particularly some specific findings, are difficult to reproduce for any of the synthesizers. Given these results, we advocate for a new class of mechanisms that can reorder the priorities for DP data synthesis: favor stronger guarantees for utility (as measured

by epistemic parity) and offer privacy protection with a focus on application-specific threat models and risk-assessment.

PVLDB Reference Format:

Lucas Rosenblatt, Bernease Herman, Anastasia Holovenko, Wonkwon Lee, Joshua Loftus, Elizabeth McKinnie, Taras Rumezhak, Andrii Stadnik, Bill Howe, and Julia Stoyanovich. Epistemic Parity: Reproducibility as an Evaluation Metric for Differential Privacy. PVLDB, 16(11): 3178 - 3191, 2023. doi:10.14778/3611479.3611517

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/DataResponsibly/SynRD>.

1 INTRODUCTION

Differential privacy (DP) has been studied intensely for over a decade, and has recently enjoyed uptake in both the private and public sectors. In situations where the downstream analysis is known, one can design specialized mechanisms with high utility [41, 42]. But an active research area is to design general DP data synthesizers (henceforth, synthesizers) that model the entire data distribution, inject noise, then sample the noisy model to generate synthetic datasets intended to be broadly usable in a variety of unanticipated applications. Evidence to support claims of general utility is typically presented as results on proxy tasks over common public datasets (e.g., the ubiquitous Adult dataset [37]). Proxy tasks may include descriptive statistics, queries involving one or two variables [28, 29, 61, 62], classification accuracy [16, 61, 67], and information theoretic measures [67]. Although these proxy tasks

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 16, No. 11 ISSN 2150-8097.
doi:10.14778/3611479.3611517

*Rosenblatt is the first author, Howe and Stoyanovich are the senior authors, other authors are listed alphabetically.


are procedurally representative of real tasks, the implicit claim of generalization to practice is rarely explored.

Limited empirical evidence on relevant tasks undermines trust in the practical use of DP. The US Census Bureau adopted DP for disclosure avoidance in the 2020 census, interpreting federal law (the Census Act, 13 U.S.C. § 214, and the Confidential Information Protection and Statistical Efficiency Act of 2002) as a mandate to use advanced methods to protect against computational reconstruction attacks unforeseen when the laws were passed. But the adoption of DP for the Census was met with resistance among many in the research community, who contend that data infused with DP noise affects demographic totals [58] and exacerbates underrepresentation of minorities [25, 36]. Besides the research implications, there are potential consequences for policy: Block grants are allocated based on minority populations as measured by the census data, and underrepresentation can lead to underfunding integral services including Medicaid, Head Start, SNAP, Section 8 Housing vouchers, Pell Grants, and more [12]. Although the Census Bureau held workshops, released demonstration datasets, and published technical reports to support the community, these outreach efforts realized limited success; multiple lawsuits are still pending as of May 2023.

Despite these challenges, DP still offers stronger guarantees of disclosure protection than, and similar utility to, alternative proposals (e.g., k -anonymity, swapping [12]). DP, when used correctly, ensures that any inferences conducted on data do not reveal whether a single individual’s information (including, for example, their gender or race) was included in the data for analysis [19]. DP can therefore not only protect privacy, but also enable access to protected demographic attributes necessary for research on fairness and equity in machine learning [33].

Characterizing DP Error. A practical method of operationalizing DP is to learn a (noise-infused) model of a dataset, then sample that privatized model to generate synthetic data that can be released publicly [20, 26, 41, 55, 63, 65]. Ideally, this approach would provide a drop-in replacement for the original data that can be used in *any* downstream context to produce reasonably faithful results with strong privacy guarantees. But this ideal is unrealizable, both theoretically and practically. Overly accurate estimates of too many statistics are blatantly non-private, affording full reconstruction of the original dataset [17]. For any DP synthetic dataset, some statistics will tend to be faithful to the original data, while others will incur essentially arbitrary error. If the privacy budget is allocated uniformly across features, descriptive statistics of each individual feature will be faithful, but the conditional probabilities and marginals needed to construct the joint probability distribution, which is needed for general inference, will be unreliably noisy, and vice versa. Utility loss may also be non-uniform across subsets of a dataset, in some cases exacerbating inequity and leading to underrepresentation [4, 36] or to error rate disparities [53]. Designers of DP synthesizers must therefore make some kind of educated guess about which tasks should be preserved and which can be ignored. Further, the error introduced by DP methods can and should be incorporated into statistical models explicitly, just as other sources of error are modeled explicitly. However, current DP synthesizers tend not to provide formal descriptions of the error they introduce; this lack of error guarantees is a major drawback of private data

Figure 1: A visual finding from Fairman et al. [23], describing the rate of drug use across demographic groups (top) and our qualitative reproduction under DP, using MST at $\epsilon = e$. Agreement is subjectively high, though imperfect.



release. Our work does not address this limitation, but does help provide an empirical motivation for doing so.

Methodology. We propose an evaluation methodology for DP synthesizers based on reproducibility: that *published findings on the original dataset should be replicable on a noise-infused dataset*. We identify conclusions in the text of published papers, extract relevant findings supporting those conclusions, implement the corresponding statistical tests using the authors’ data, generate synthetic datasets using state-of-the-art DP synthesizers, re-apply the statistical tests over the synthetic data, and then determine if the findings still hold. If all findings hold, we say that the DP synthesizer achieves *epistemic parity* for that paper.

We instantiate our methodology over a benchmark of peer-reviewed sociology papers that are based on public data from the Inter-university Consortium for Political and Social Research (ICPSR) repository. We model quantitative results as an inequality between two numbers, for example, “Those using marijuana first (vs. alcohol or cigarettes first) were more likely to be Black, American Indian/Alaskan Native, multiracial, or Hispanic than White or Asian.” [23] In addition to quantitative results, we consider qualitative results: a subjective assessment of whether key visualizations in the paper expose the same relationships when recreated over synthesized data. Figure 1 gives an example, presenting a visualization from Fairman et al. [23]. In this figure, we show the break-down of drug use across demographic groups, taken directly from the publication (top of the figure), and compare it with a reproduction of the plot based on DP synthetic data generated with the MST synthesizer at $\epsilon = e$ (bottom of the figure).

Following Errington et al. [22], and as is common in the reproducibility literature, our aim was not to reproduce every finding in

every paper; rather it was to identify and reproduce a selection of key findings from each paper. For generality, interpretability and simplicity, we consider whether a conclusion holds over synthetic data to be true if the the two quantities are in the same relative order, and do not attempt to measure the change in effect size or the statistical significance of the difference between the original and synthetic result.

Benchmark and results. ICPSR is an NSF-funded repository for social science data holding over 100,000 publications associated with 17,312 studies. A study typically involves hundreds of variables and supports dozens of papers. Each paper can be considered to be deriving its own dataset (selected variables and selected rows) from the source data of the study. We apply DP methods to synthesize data for these paper-specific, study-derived datasets.

ICPSR studies are publicly available by policy, which enables us to instantiate the epistemic parity methodology and develop a benchmark. Notably, there is increasing demand from the ICPSR leadership and community to support keeping sensitive data private, while generating DP synthetic subsets to support reproducibility. Our methodology can be used to respond to this demand.

Paper selection. The benchmark we define consists of 4 datasets and eight recent peer-reviewed papers selected for impact, accessibility of the topic to non-experts, recency, and several other criteria. We extracted findings and attempted to reproduce them, following the "same data, different code, different team" approach to reproducibility, encountering challenges commonly reported in that literature including undocumented data versioning, unspecified or incomplete methodologies, and irreconcilable differences between our reproduction and what the authors report. Each paper received, at a minimum, attention from two researchers with advanced degrees in either computer science, statistics, or both, and at least thirty hours of work. A complete list of papers that we attempted to reproduce, and the issues we encountered during reproduction, in our public GitHub repository, under archive, available at <https://github.com/DataResponsibly/SynRD>.

DP synthesizer selection. We use six state-of-the-art DP synthesizers, namely, MST [41], PrivBayes [67], PATECTGAN [55], AIM [42], PrivMRF [10], and GEM [39] executing each at their recommended settings. We describe these methods in Section 2.

Summary of results. We find that marginals-based and Bayes-net based state-of-the-art DP synthesizers are able to achieve high epistemic parity for five out of eight papers in our benchmark, but that some papers, and particularly some specific findings, are difficult to reproduce for any of the synthesizers, suggesting a basis for a new benchmark. The papers on which high epistemic parity is achieved use relatively low-dimensional tabular data. However, as we show empirically, large domain and high-dimensional settings are still a bottleneck for increased adoption of DP synthesizers.

Roadmap and Contributions. We discuss related work on reproducibility, DP synthesis and evaluation of DP in Section 2, give background on DP in Section 3, and then present our contributions:

- We propose the epistemic parity evaluation methodology, based on reproducing qualitative and quantitative empirical findings in peer-reviewed papers over DP synthetic datasets (Section 4).

- We instantiate the epistemic parity methodology for eight peer-reviewed social science publications, creating a reusable benchmark for evaluating synthesizers (Section 5).
- We present SynRD, an open-source benchmarking package that automates epistemic parity evaluation, and can be extended to include additional DP synthesizers, publications, and finding types (Section 6).
- We present experimental results on our benchmark, using five state-of-the-art DP synthesizers (Section 7).

We conclude with a discussion of the results, identifying trade-offs and motivating a new class of privacy techniques that favor strong epistemic parity and de-emphasize privacy risk, in Section 8.

2 RELATED WORK

DP synthesis. In our evaluation, we considered five state-of-the-art private data release methods: MST, AIM, PrivMRF, PATECTGAN, PrivBayes and GEM. We acknowledge that many other methods exist for generating DP data [20, 26, 63, 65]. We chose this set informed by recent work by McKenna et al. [42], Tao et al. [62] showing that, over randomized query workloads on tabular data, MST, AIM and PrivMRF are the highest-performing marginal-based methods, that PrivBayes is the highest-performing Bayes-net-based method, and that PATECTGAN and GEM are the highest-performing deep learning based methods. AIM, PrivMRF and GEM are more recent than MST; they were not included in the recent dedicated DP synthesizer benchmarking survey by Tao et al. [62] and are currently considered to be the state-of-the-art DP synthesizers.

PrivBayes derives a Bayesian model and adds noise to all k -way correlations to ensure differential privacy [67]. This method was published in 2017, yet it remains competitive with and even outperforms other methods, described next. PrivBayes has demonstrated efficacy in a number of settings, but is computationally inefficient for high-dimensional data.

MST relies on the Private-PGM graphical model to construct a maximum spanning tree among attributes in the data feature space, where edges are weighted by mutual information [40]. By measuring 1-, 2- and some 3-way marginals, MST is able to create a high-fidelity low-dimensional approximation of the joint distribution between all attributes, leading to impressive statistical utility on metrics like mean, standard deviation, and bivariate correlations.

AIM, like MST, relies on the Private-PGM for parameterizing the underlying distribution [42]. However, unlike MST, AIM is *workload aware*: It parameterizes a private synthetic distribution through an iterative process akin to that of the fundamental DP synthesis Multiplicative Weights Exponential Mechanism (MWEM) algorithm [26], taking advantage of a set of queries of interest, specified *a priori*. In their evaluation, the authors of AIM found that it outperformed MST convincingly, especially for higher values of ϵ .

PrivMRF is another marginal-based algorithm that relies on Private-PGM [10]. Its novelty lies in a clever selection of marginals to measure, which aligns with three main criteria: marginals should be low-dimensional, the graph of marginals should be small, and any junction tree transform on the graph of marginals should not result in domain blowup. PrivMRF is competitive with PrivBayes and MST, as demonstrated by both [10] and [42].

PATECTGAN relies on a conditional generative adversarial network tuned to tabular data, where the discriminator has privacy constraints [55, 66]. The high up-front cost of initializing and updating the weights in a full discriminator network means that PATECTGAN is limited in low dimensional settings, but shows particular promise when the data is higher-dimensional.


GEM analyzes the “Adaptive Measurements” framework for private synthetic data algorithms, inspired by the MWEM architecture [26]. Iteratively, an algorithm (1) privately selects a set of queries; (2) obtains noisy measurements of these queries; and (3) updates an approximating distribution according to some loss function. The GEM algorithm uses a similar loss function to the RAP private data release algorithm [3], although with small modifications that lead to significantly improved practical performance [39, 42].

Workload-aware synthesizers. The AIM and GEM synthesizers in our study would be considered workload-aware in that they are able to adaptively select marginals based on some *a priori* specified workload of interest. In our setting, scientists pre-select a relatively small subset of variables (say, twenty) of interest for analysis from a large study with thousands of variables. We consider this pre-selection, with any relationships among them permitted, to be the workload specification; further limiting the queries would problematically bias the data outcomes. We thus assume that relationships between any of the selected variables of interest are permitted. Thus, for the AIM and GEM algorithms, we automatically generate a workload across all variables.

DP evaluation. Tao et al. [62] evaluated several DP synthesizers’ ability to preserve 1- and 2-variable distributions. Jayaraman and Evans [34] studied privacy-utility tradeoffs for ML tasks and found that commonly-used ϵ values and implementations in practice may be ineffective: either unacceptable privacy leakage or unacceptable utility tends to occur. Hay et al. [28] used 1- and 2-dimensional range queries over 27 public datasets to study the influence of dataset scale and shape that had led to inconsistent results in the previous literature. While the datasets are diverse in relevant properties, the tasks are limited, and the link to conclusions drawn is unexplored. Hill [29] studied the utility of DP on one longitudinal behavioral science dataset involving a sexuality survey motivated by real world attacks based on disclosing pregnancy, finding that theoretical guarantees of DP were generally supported, but that high-dimensional data was a challenge for utility. In this work, we aim to facilitate and standardize these kinds of applied studies.

Reproducibility. Numerous reproducibility studies have been attempted in various fields, typically reporting remarkably low rates of success, leading to calls for significant changes to policy and incentive structures underlying scientific funding and publishing [43–45]. In the social sciences, Camerer et al. [11] replicated 21 systematically selected experimental studies published in Nature and Science between 2010 and 2015, finding a significant effect in the same direction as the original study for 13 (62%) studies, with about half the effect size, on average. The Reproducibility Project: Cancer Biology Errington et al. [22] attempted to reproduce 193 experiments from 53 papers, but succeeded in only 50 experiments from 23 papers. They found that only 2% of papers supplied open data, 0% of protocols were completely described, 67% of experiments required

Figure 2: Epistemic parity workflow: Each study dataset supports many papers, each using a subset of the features. The paper’s findings are implemented as computable inequalities. We generate many privatized datasets using different random seeds, then compute the proportion of these trials for which the findings hold (Figure 3).



modifications to complete, and replication effect sizes were 85% smaller than in the original findings. A survey by Baker [5] found that 52% of respondents agreed that reproducibility represents a ‘crisis’ for science.

The terms *reproducibility* and *replicability* are used inconsistently across and even within fields [6]. We use reproducibility to mean same data, different code, different team [6]. We use replicability to mean acquiring new data in a new experiment to determine whether the same conclusions hold. Our methodology amounts to first *reproducing* conclusions found in peer-reviewed papers, and then *replicating* these same conclusions on DP synthetic data. While our focus is neither to reproduce nor to replicate the original study, our framework supports reproducibility analysis as an intermediate step, and our insights regarding the difficulty (and, often, the impossibility) of reproduction are consistent with prior findings.

3 PRELIMINARIES

Differential privacy (DP) is a strong privacy guarantee that ensures that altering or removing one record from a given dataset does not significantly affect the outcome of an analysis or query. Intuitively, DP prevents an observer of a private output from drawing conclusions about which specific individuals’ information was included in the input. DP is based on the concept of “neighboring” datasets, which are datasets that result from the modification or removal of one record that induces a “neighboring” dataset. There is a distinction between modification and removal in this definition; establishing which version of the privacy definition is normally done *a priori*. In the scope of the private synthesizers considered by this paper, datasets X and X' are considered neighboring if the removal of a single element x_i from one yields the other (except in the case of PrivBayes; we account for this in our budget allocation).

The conventional understanding of DP is provided in Definition 3, while the alternative, zero-concentrated DP (ρ -zCDP) is provided in Definition 3.

Definition ((ϵ, δ) -Differential Privacy). A randomized mechanism \mathcal{M} is said to be (ϵ, δ) -differentially private if, for every pair of neighboring datasets X and X' , and all subsets R of possible outputs, the following holds:

$$\Pr[\mathcal{M}(X) \in R] \leq e^\epsilon \Pr[\mathcal{M}(X') \in R] + \delta$$

Definition (Concentrated Differential Privacy (ρ -zCDP)[9]). We use $D_\alpha(\mathcal{M}(X)||\mathcal{M}(X'))$ to signify the α -Rényi divergence. Thus, a randomized algorithm \mathcal{M} upholds ρ -zCDP if, for every pair of neighboring datasets X and X' ,

$$D_\alpha(\mathcal{M}(X)||\mathcal{M}(X')) \leq \rho\alpha, \forall \alpha \in (1, \infty)$$

These two principles, (ϵ, δ) -DP and ρ -zCDP, are intimately related but scale with different relative privacy parameters. As demonstrated by Bunet *et al.* [9], an established hierarchy of these guarantees is: An $(\epsilon, 0)$ -DP mechanism gives $\frac{\epsilon^2}{2}$ -zCDP, which in turn gives $(\epsilon\sqrt{2\log(1/\delta)}, \delta)$ -DP for every $\delta > 0$.

It is theoretically difficult to compare private mechanisms satisfying differing versions of DP. However, in the case of private synthesizers, practical, utility-based comparisons have become a norm [10, 42, 55, 62]. The methods we consider provide different guarantees: AIM and GEM both give ρ -zCDP guarantees, while MST, PATECTGAN and PrivMRF give (ϵ, δ) -DP guarantees, and PrivBayes gives an $(\epsilon, 0)$ -DP guarantee. For comparison purposes, all ϵ parameters are translated using the parameter relationships given above, so as to compare at the same relative privacy settings between synthesizers. As is accepted practice, we set our δ parameter to be “cryptographically small,” at the very most $\frac{1}{n}$ where n is the size of our data, but usually much smaller [21].

4 THE EPISTEMIC PARITY METHODOLOGY

Informally, we say that epistemic parity holds if all conclusions one draws on the *original dataset* also hold on the *synthetic dataset*. We define epistemic parity operationally as follows.

Definition (Epistemic Parity). Let D be a universe of datasets. A *finding* is a function $D \rightarrow \{0, 1\}$. A *synthesizer* is a function $S : (X, D) \rightarrow D$ where $X \sim U(0, 1)$ is a random seed. Given a dataset $d \in D$, an associated finding f , and a DP synthesizer S , we say the epistemic parity of S for f on d is the fraction of random trials $i \in X$ such that $f(S(i, D)) = f(D)$.

The overall workflow to compute epistemic parity is illustrated in Figure 2. The input is a set of papers, and the output is a set of scores indicating whether findings are supported under various DP synthesizers. This framework, and the associated open source software that partially automates it, is intended to expand over time. Mechanism designers are encouraged to implement new DP methods, and researchers are encouraged to implement new findings from relevant papers.

Beginning from the left in Figure 2: A study is associated with one dataset and potentially many papers. Each paper typically uses only a subset of the variables in the study. We assume public access to the data on which the paper’s results were computed; our focus is on evaluating DP methods (requiring ground truth) rather than on protecting the privacy of subjects involved in the study. The vast majority of evaluations of DP mechanisms requires access to the original data for interpretation; indeed, inaccessible ground truth undermined the US Census Bureau’s efforts to build trust regarding DP [8].

Given a paper, we manually identify natural language claims made by the authors as candidates for findings. Though these claims may appear anywhere in the paper, most were found in the results

section. Domain expertise of course provides an advantage in this exercise, but we contend that it should always be possible for non-expert readers to identify major claims since the goal of a paper is to communicate findings to a broader audience.

For each claim, we identify the quantitative (or qualitative) evidence that supports the claim, recording the variables involved and methods used. We then re-implement the analysis to (attempt to) reproduce the salient findings and conclusions in the paper over the original, public dataset, as discussed in Section 4.1.

While in principle this reproduction step is always possible, in practice, it can be difficult or impossible [5, 44], and may involve guesswork when the computational details are incomplete. Moreover, inconsistent reproducibility can introduce bias in our benchmark: we may be more likely to include findings for which computational details are clear, which may be those that are simpler to explain or better-known by the author, as opposed to a representative sample of appropriate methods.

If the reproduction was successful, we generate $k \times m$ synthetic datasets representing k trials with different random seeds and m different DP methods, and then draw an additional B samples from each seeded DP method. In our initial benchmark, $k = 10$ and $m = 5$, and $B = 25$. The additional B draws allow us to bootstrap a confidence interval for each (trained) synthesizer. That is, there are two sources of randomness: the training procedure used by the mechanism, and the random sampling of the learned model to actually generate synthetic data. Although each synthetic dataset could be scaled to any number of records — recall that we are sampling a privatized model — we always use the same number of records as the original data for each bootstrap sample. Given this set of synthetic datasets, we again attempt to reproduce the findings using each one. Finally, we contrast the findings based on original and DP data by measuring the proportion of trials, for each method, where a given finding holds. Our methodology is implemented in an open-source framework, described in Section 6.

4.1 Reproducing Experimental Studies

We adapt three concepts of reproducibility—*values*, *findings*, and *conclusions*—from Cohen *et al.* [13] into a practical taxonomy for reproducing a statistical analysis in a peer-reviewed publication, and implement a software framework that allows us to conduct concrete experiments around this taxonomy (see Section 6).

The atomic element in reproducibility is a *finding*, defined by Cohen *et al.* [13] as “a relationship between the *values* for some reported figure of merit with respect to two or more dependent variables.” For the purposes of our study, a *finding* consists of a natural language statement (i.e., a *claim*) reported in a publication, along with evidence provided by one or more quantitative or qualitative sub-statements about the analysis.

Evidence for a *finding* consists of a comparison between two or more *values* that can be evaluated as a Boolean condition. A value may be a scalar (i.e., 34.1%), an aggregated or computed result (i.e., a regression coefficient of 1.2), or even an implicit threshold expressed in natural language (e.g., “a low rate” or “a strong correlation”). In these cases, we instantiate the language as a quantitative threshold, applying conventions from the literature when they exist.

For example, a common convention is that Pearson’s correlation is considered “strong” when r is larger than 0.7.

A special case of a *finding* is a qualitative *visual finding* that often appears in the form of a figure, table or diagram. A figure encodes many potential *findings*; we do not (necessarily) consider each of these sub-findings on their own in our analysis, but rather treat them as a single *visual finding*: we attempt to reproduce the figure itself, and subjectively evaluate its similarity to the original. Consider Figure 1, where the top sub-figure shows a percentage breakdown of drug use across demographic groups from Fairman et al. [23], and the bottom sub-figure shows the same breakdown reproduced on DP synthetic data generated by MST at $\epsilon = e$.

Finally, following Cohen et al. [13], a *conclusion* is defined as “a broad induction that is made based on the results of the reported research.” A conclusion must be explicitly stated in a paper, and comprises one or several *findings*. Conclusions are less concrete than individual findings, and are often at the discretion of the domain expert tasked with interpreting a set of findings.

4.2 Generating DP Synthetic Data

Each of the papers that we reproduced using DP synthetic data derived findings from a subset of the full study’s data. For example, HSLs:09 consists of over 7000 columns, but Jeong et al. [35] used only a subset of 57. We synthesize the subset of data relevant for the reproduced findings and conclusions, as discussed in Section 4.1. In the case where a paper relies on longitudinal data from a study, we collapse the data such that it is “one row to one person.”

The DP methods for private data release are executed for the range of ϵ values $\epsilon \in \{e^{-3}, e^{-2}, e^{-1}, e^0, e^1, e^2\}$, which represents a small to medium privacy regime [7]. Each DP mechanism is run 10 times to produce, at each ϵ value, $10 \times B$ sampled datasets using the same sample size but different random seeds (where B is the bootstrap parameter). Each DP method involves different hyperparameters and varying levels of tunability, but we use author-recommended settings to avoid biasing results towards our own expertise. We then re-compute the findings for each sample.

If all findings are reproduced regardless of *epsilon* or random seed, we say that the DP mechanism achieves complete epistemic parity. But we measure parity as the proportion of iterations for which the finding holds. The goal is to overlook small variations in the exact value in favor of maintaining the relative relationships of the computed statistics for interpretability and practical utility.

4.3 Analysis of Variability for DP Synthetic Data

Räisä et al. [51] considers a set of rules proposed by Raghunathan and Rubin as a potential approach to calculating uncertainty over the results of DP synthetic data [50]. Consider synthetic datasets $x_1, \dots, x_m \in X$, and test statistic τ . Running $\tau(X_i)$ produces a point estimate q_i with variance v_i for dataset X_i . Standard hypothesis testing might calculate a confidence interval and power over the distribution of all q_i . In the case where X consists of synthetic data, Rubin’s rules for producing such calculations can be applied under normality assumptions [50, 51]. In brief, these rules are as follows: to compute confidence intervals, the estimated locations q_1, \dots, q_m

and variances v_1, \dots, v_m are combined:

$$\hat{q} = \frac{1}{m} \sum_{i=1}^m q_i \quad \hat{v} = \frac{1}{m} \sum_{i=1}^m v_i \quad (1)$$

$$b = \frac{1}{m-1} \sum_{i=1}^m (q_i - \hat{q})^2 \quad (2)$$

And then we use a t -distribution centered at \hat{q} , with variance

$$T = \left(1 + \frac{1}{m}\right) b - \hat{v}, \quad (3)$$

and degrees of freedom:

$$\text{df} = \left(1 - \frac{1}{1 + \frac{1}{m} \frac{\hat{v}}{b}}\right)^2 (m-1). \quad (4)$$

This method *crucially* relies on a normality assumption for the distribution over each $\tau(X_i) = q_i$. If the τ statistic is a mean and variance over a well understood distribution, this could be a reasonable assumption to make. However, in the case of our benchmark, we consider a wide variety of more complex, multivariate statistical tests over multiple high dimensional datasets. For example, weighted least-squares regression produces a set of coefficients for covariates, and we compare the magnitude of these covariates in relation to a target variable of interest. Assuming a normal distribution over covariate values across different randomly initiated trainings of a synthesizer, with the addition of DP noise, is strong. Thus, we are uncomfortable making this normality assumption.

We take a more conservative approach. Instead of computing the statistical power of each point estimate q_i , we turn each test into a “soft” finding, given informally by Equation 5. The function *finding* here takes in τ , some ground truth data q_i^{real} and synthetic data q_i^{synth} drawn from a privatized distribution intending to match the distribution of q_i^{real} . The parameter α manages the tolerance of the comparison, acknowledging that the DP noise and the nature of synthetic data will certainly alter the true values slightly. α is necessarily assigned according to the context of a *finding* (for example, if for groups A and B, we compare $\mu_A, \mu_B \in (0, 1)$, where in the real data $\mu_A \approx \mu_B$, a tolerance of $\alpha = 0.01$ might suffice).

$$\text{finding}(\tau, q_i^{synth}, q_i^{real}) = \mathbb{1}[|\tau(q_i^{synth}) - \tau(q_i^{real})| \leq \alpha] \quad (5)$$

Recall that the goal of our benchmark is primarily to determine the viability of a synthesizer for reproducing a specific paper’s findings, so that we draw the same overall conclusions. Parity is then the proportion of agreement for a given paper and given synthesizer over a fixed set of discrete, statistical tests posed as *findings*. This approach allows us to compare across findings of different types, and to score synthesizers for each paper on the reproducibility of their results.

But how should we address the multiple sources of *variance* in synthetic data? There are two main elements of randomness. Each parametric DP synthesizer is trained conditioned on a random seed, which dictates the privatizing random noise introduced during model fitting. Varying the seed changes the random initialization bias *significantly* for DP synthesizers. Then, after training, once a synthesizer is parameterized, a specific draw from that synthetic distribution is again conditioned on a *different* random seed, which

dictates the bias of a particular realized sample from the synthetic distribution. The variance effect for the post training draw is much lower, as no private noise need be added during sampling.

The latter source of randomness (in drawing from a synthetic distribution) is easier to deal with. Conditioned on synthesizer S_i , for fixed finding $finding(\tau, q_i^{synth}, q_i^{real})$, we perform a bootstrap over B samples of size n (where n is the size of the real ground truth empirical sample we trained on). Over B datasets, we calculate the average statistic and average variance of $finding$, $\hat{q} = \frac{1}{m} \sum^m q_i$ and $\hat{v} = \frac{1}{m} \sum^m v_i$, as well as a 95% confidence interval.

The former source of randomness (in fitting a synthetic distribution) is more difficult to deal with, as fitting a synthetic distribution is computationally expensive for high dimensional data [42, 55, 62]. Despite running extensive tests, we were unable to run enough trainings per synthesizer to say anything convincing about the statistical power of our sample of random synthesizer initializations. Instead, we simply report on the variance of the aforementioned soft finding bootstraps over 10 randomly initialized synthesizers, per paper, per privacy level (around 1400 trained synthesizers in total). Though this is standard practice [4, 7, 62], we believe there’s an opportunity to improve the norms here through model sharing and well supported open source synthesizers.

5 BENCHMARK CONSTRUCTION

We used a standardized approach for study and paper selection. Each study, which has an associated dataset, was selected for broad impact (at least 100 papers). For each selected study, we selected peer-reviewed papers published during the past 5 years that are publicly available, use publicly available data, and are no more than 30 pages long. We describe selected studies and papers below, see <https://github.com/DataResponsibly/SynRD> and the full version of the paper [54] for details.

5.1 Selected Studies

HLSL:09: High School Longitudinal Study [14], is a nationally representative, longitudinal study of U.S. 9th graders who were followed through their secondary and postsecondary years. We attempted to reproduce four papers that use HLSL:09, and were able to fully or partially reproduce three of them.

AddHealth, National Longitudinal Study of Adolescent and Adult Health [27], consists of a nationally representative sample of U.S. adolescents in grades 7 through 12 during the 1994-1995 school year. We attempted to reproduce four papers that use Add Health, and were able to fully or partially reproduce two of them.

NSDUH, National Survey on Drug Use and Health 2004-2014 [64], measures the prevalence and correlates of drug use in the U.S.. We attempted to reproduce four papers that use NSDUH, but we were able to reproduce only one, partially and with substantial effort. The main obstacle was that the study is broken down by year, and multiple versions of the study for each year have been released, without a record of what the new version modifies.

ACL, The Americans’ Changing Lives Survey [30], is an ongoing longitudinal study of the lives of U.S. adults. The study has several waves, the first of which was conducted in 1986, and each wave follows up with the same sample of U.S. adults, asking questions to determine how social connections, work, and other factors affect

health throughout their lifetimes. We attempted to reproduce 2 papers that use ACL, and were able to partially reproduce both.

5.2 Selected Papers

Saw et al. [59] use HLSL:09 to examine cross-sectional and longitudinal disparities in STEM career aspirations among high school students. Methods include singular and trivariate disparity analysis w.r.t. gender, race/ethnicity and socioeconomic status, and analysis of disparities among students who persist in their STEM interest from 9th to 11th grade vs. those who first emerge with STEM interest in 11th grade. This paper was reproducible, with effort - we were able to replicate each finding and agree with all conclusions.

Lee and Simpkins [38] use HLSL:09 to evaluate the impact of factors, such as “low teacher support,” and self-perceptions of “math ability” and “parental support”, on the performance of students on the 11th grade math exam. The authors construct these factors by aggregating across survey responses with a weighted average, and control for demographic variables and historical math performance to isolate the effects. They compute Pearson’s correlation and train linear regression models to predict math performance. This paper was partially reproducible, with substantial effort. The authors did not detail their regression techniques, although results of a simple weighted linear regression turned out to align well with their findings. We were unable to reproduce a finding that involves a covariance slope analysis figure from a complex R package.

Jeong et al. [35] use HLSL:09 to interrogate potential racial bias in classifiers that predict student performance on a standardized 9th grade math exam. The authors assign each student the privileged (White/Asian) or disadvantaged (Black/Hispanic/Native American) group. They trained random forest, SVM and logistic regression models to predict whether a student would receive a top-50% test score. We found this paper to be reproducible, with effort. The authors did not specify how the data was preprocessed (e.g., how missing values were imputed) or how it was split into training and test. We were ultimately able to reproduce the results sufficiently to agree with the conclusions, but were unable to reproduce the values in the findings exactly.

Fruitt and Chan [24] use AddHealth to investigate the role of mentors in the educational outcomes of first-generation college students. The authors fit a statistical mediation model by Preacher and Hayes [49] (PROCESS Model, variation 1) to test direct and interaction effects of parental educational attainment and mentorship on students’ educational attainment. This paper was partially reproducible, with substantial effort. We were unable to reproduce findings based on manual qualitative coding of free text responses, because the authors did not publish the coding scheme. Still, we were able to reproduce most of the observed trends and agree with the conclusions.

Iverson and Terry [31] analyze the effect of having played high school football on depressive and suicidal tendencies in men later in life using the first wave (1994-95) and the most recent wave (2016-2018) of AddHealth. The authors conduct a bivariate analysis across two groups of men (those who did or did not play football in high school), controlling for demographics and risk factors, and report on simple percent comparisons, statistical significance and odds ratios. This paper was easy to reproduce. The authors provided precise

Table 1: Properties and meta-features of the datasets in our benchmark, and of two datasets that are commonly used for DP benchmarking, Adult and Mushroom. Mutual Information, Skewness and Sparsity are the *average* for each of these metrics across all variables in the dataset. Our results (Section 7.1) reinforce that synthesizers may struggle with large sample sizes (Fairman et al. [23]), large domain sizes (Jeong et al. [35]), and low mutual information (Iverson and Terry [31]).

Paper	Sample Size	Variables	Domain Size	Outliers	Mutual Info.	Skewness	Sparsity
Assari and Bazargan [2]	3361	16	9.03e+09	9	0.051 ± 0.153	0.563 ± 1.557	0.253 ± 0.231
Fairman et al. [23]	293581	6	2.03e+05	0	0.255 ± 0.432	0.185 ± 0.462	0.174 ± 0.165
Fruiht and Chan [24]	4173	11	2.20e+05	6	0.104 ± 0.256	0.607 ± 1.694	0.394 ± 0.183
Iverson and Terry [31]	1762	27	5.71e+15	5	0.004 ± 0.010	NaN	0.307 ± 0.180
Jeong et al. [35]	15054	57	7.04e+42	32	0.020 ± 0.026	0.338 ± 2.850	0.261 ± 0.166
Lee and Simpkins [38]	14575	9	5.11e+17	5	2.862 ± 1.242	0.080 ± 0.440	0.111 ± 0.156
Pierce and Quiroz [46]	1585	17	7.19e+11	11	0.030 ± 0.050	0.001 ± 1.050	0.146 ± 0.158
Saw et al. [59]	20242	9	4.30e+04	3	0.143 ± 0.145	1.291 ± 1.218	0.354 ± 0.171
Adult [37]	32561	15	9.06e+14	96	0.066 ± 0.053	17.455 ± 22.992	0.125 ± 0.164
Mushroom [60]	8124	23	2.44e+14	74	0.199 ± 0.209	6.211 ± 8.955	0.297 ± 0.219

guidance on which questions (columns) they used and applied relatively simple methods. We were able to exactly match nearly all values reported in the paper, reproducing all findings.

Fairman et al. [23] use NSDUH to investigate predictors and potential consequences of initiating the use of marijuana before other types of substance (e.g., cigarettes and alcohol) for U.S. youth. The primary methods of analysis were counts and percentage comparisons by group, and computing adjusted relative risk ratio (aRRR) and adjusted odds ratio (aOR). This paper was partially reproducible, with substantial effort. The main obstacle to reproducibility is the inappropriate versioning of the data by NSDUH, as discussed in Section 5.1. Additionally, the paper inadequately describes their calculations of aRRR and aOR, making them difficult to replicate. We were unable to reproduce exact reported values but reproduced general trends and agree with most conclusions.

Assari and Bazargan [2] use ACL to investigate the predictive role of obesity on long-term risk of mortality due to cerebrovascular disease, and, further, to test racial variation in this effect. The authors use the first wave of ACL (1986) to establish baseline factors, including obesity as the main predictor, and demographics, socioeconomic status, health behaviors, and health status as baseline covariates. They used data from subsequent waves to determine time to death due to cerebrovascular causes. This paper was partially reproducible, with substantial effort. To reverse-engineer their analysis, we used history of cerebrovascular disease combined with death to represent this feature, and successfully reproduced 16 of the 18 findings we attempted.

Pierce and Quiroz [46] use ACL to study how social support and social strain stemming from spouses, children, and friends impact emotional states. The authors modeled the variables by combining answers from multiple survey questions and used a mixed-effects model. This paper was partially reproducible, with substantial effort. The authors did not clearly state which variables they used as basis for demographic attributes, and we were thus unable to fully reproduce their analysis. We were able to reproduce 10 out of 14 reported findings.

Note on study/dataset dimensionality. We did not explicitly filter papers based on the size of the dataset they used. The studies we considered were very high dimensional (many thousands of

variables), but the corresponding papers in our benchmark each follow a standard subsetting procedure, where they select a small collection of variables of interest for analysis. Thus, our benchmark datasets are not as high-dimensional as other benchmarks [42].

5.3 Comparison to Other DP Benchmarks

Our benchmark papers introduce a collection of 8 new datasets. In this section, we briefly analyze these new datasets, and discuss characteristics that differentiate them from standard ML datasets, like Adult [37] or Mushroom [60], which were used in prior DP benchmarks [28, 55]. We perform this analysis through the lens of principled dataset and ML task characterization, sometimes called “meta-learning.” Prior work in the meta-learning community provides guidance [48, 52] and tooling [1], of which we take advantage.

In Table 1, we show several properties and meta-features for eight datasets from our benchmark, as well as for two popular datasets from the UCI Machine Learning repository [18], Adult [37] and Mushroom [60]. Four of the meta features (*Number of outliers*, *Mutual information*, *Skewness*, and *Sparsity*) might be of particular interest to DP data synthesis. We discuss these further in the full version of the paper [54].

Overall, Table 1 demonstrates that the datasets in our benchmark are diverse and cover a wide range of values for the aforementioned metrics. Interestingly, one of our most challenging datasets to reproduce, Iverson and Terry [31], had the lowest average mutual information score and one of the highest sparsity scores. Many of the synthesizers we test depend on mutual information to select the marginal measurements for distribution learning. Selecting the most relevant 2-way marginals when mutual information is uniformly *low* and there are many features is clearly a challenge. We further note that Adult, a common challenge dataset, had uniquely skewed distributions, which aligns with prior work suggesting that this dataset is idiosyncratic, and so one should reconsider using it for evaluation and benchmarking [15].

6 EXTENSIBLE BENCHMARKING PACKAGE

We organize the epistemic parity workflow and existing papers and datasets into an open-source differentially private (DP) synthetic data benchmarking package called SynRD, available at <https://>

`//github.com/DataResponsibly/SynRD`. This package helps answer the question: “Can a DP synthesizer produce private (tabular) data that preserves scientific findings?” In other words, to what degree does a DP synthesizer satisfy epistemic parity for a particular set of findings?

We used SynRD to automate empirical evaluation of epistemic parity of five state-of-the-art DP synthesizers (see Section 2) over all publications in our benchmark (see Section 5). The SynRD package is built to be naturally extensible to include additional DP synthesizers, datasets, publications, and finding types.

The main classes implemented by SynRD are Synthesizer, Publication, Finding, and Benchmark. The Synthesizer class provides a unified interface to the implementations of five DP synthesizers, specifying recommended parameter values for each, and implementing the `fit` and `sample` methods. This class wraps implementations of MST, PATECTGAN and AIM from the SmartNoise package, and an implementation of PrivBayes from the DataSynthesizer package [47]. PrivMRF is implemented separately.

The Publication class contains descriptive attributes of each paper, specifies input files and any necessary data transformations. This class also contains a collection of findings, described next. The Finding class implements each analysis method from Table 2 as finding type, and automates the process of checking whether a finding is reproduced over a given dataset.

Finally, the Benchmark class takes as input a publication, along with execution parameters (such as the number of bootstrap samples), checks all findings for that publication over the real data, generates synthetic datasets for each DP synthesizer, checks findings over synthetic data, and, finally, generates an epistemic parity score for each (synthesizer, finding) pair, and for the synthesizer overall (over all findings).

An example instance of Publication, for Saw et al. [59], along with an implementation of evaluating epistemic parity for MST using Benchmark, is shown below.

```
from SynRD.papers import Saw2018Cross
from SynRD.benchmark import Benchmark
from SynRD.synthesizers import MSTSynthesizer
benchmark = Benchmark()
B = 25 # Bootstrap parameter
synth = MSTSynthesizer(epsilon=1.0)
papers = benchmark.initialize_papers([Saw2018Cross])
for paper in papers:
    synth.fit(paper.real_dataframe)
    dataset = synth.sample(len(paper.real_dataframe) * B)
    paper.set_synthetic_dataframe(dataset)
    benchmark.eval(paper, B=B)
```

An example function that implements checking a finding that “Obesity at baseline was not associated with cerebrovascular death in the pooled sample” from Assari and Bazargan [2], is shown below.

```
def finding_6_9(self):
    corr_df = self.get_corr()
    corr_obesity_death = \
        corr_df['Obesity'].loc['Cerebrovascular death']
    soft_finding = abs(corr_obesity_death) < 0.05
```

To support empirical evaluation of our benchmark, we pre-processed publication-relevant subsets of the ICPSR studies, and deposited them in the Harvard Dataverse repository. The auxiliary DatasetLoader class retrieves this data, and makes it available to an instance of the Publication class.

7 RESULTS

Computational resources. Our benchmark consists of eight papers, each evaluated on six synthetic data algorithms for six values of ϵ , for a total of 36 mechanisms for each paper, each repeated with 10 random seeds. We draw 25 samples of size n , where n is the real data sample size, and bootstrap over this set of samples when calculating average parity over our finding set. Benchmarking extensively with DP synthesizers is computationally expensive [42, 55, 62]. Fitting many synthesizers took 100s of compute hours. Training PrivMRF and PATECTGAN was done using NYU’s Greene High Performance Computing cluster using A100 and RTX8000 NVIDIA GPUs with 80GB and 48GB of RAM respectively. CPUs from that same cluster were used to train AIM, MST, PrivBayes, and GEM. The benchmark itself (assessing parity per paper) was also run on the cluster. Though this was computationally less intense (with some exceptions), bootstrapping to report variance per synthesizer did incur additional overhead.

7.1 Quantitative results


Figure 3 shows parity for all findings across all papers, for each of the five synthesizers, with ϵ regimens of e^{-3} , e^{-2} , e^{-1} , e^0 , e^1 , and e^2 . Darker color indicates lower average parity, while lighter indicates higher average parity. Each paper is a block of rectangles, where the x -axis represents *findings* and the y -axis shows the five synthesizers. The crosshatched cells indicate that a synthesizer was unable to fit to a dataset in under 6 hours.

Baseline: Bayesian bootstrap. An important element of our evaluation framework is an accurate assessment of replicability and variability. We subject private synthesizers to a rigorous bootstrap process to assess the epistemic parity of findings presented in each paper; however, it is then reasonable to ask whether the findings themselves, assessed via a bootstrap procedure on the *real* data, would hold up under similar scrutiny? In order to check this, we perform a standard Bayesian bootstrap procedure [57]. The Bayesian bootstrap assigns weights to each observation in a given dataset, where those weights $\mathbf{w} = (w_1, w_2, \dots, w_n)$ are drawn from a symmetric Dirichlet distribution, i.e. $\mathbf{w} \sim \mathcal{D}(1, 1, \dots, 1)$ (the Dirichlet distribution is a standard Bayesian prior, a smoothed version of the multinomial distribution). We then compute the epistemic parity of each finding with these new weighted samples for B iterations of the bootstrap, matching the number of iterations of the bootstrapping procedure on the private synthetic data. Each statistic is then an independent and identically distributed realization from the posterior predictive distribution, which offers a Bayesian perspective on the data-generating mechanism. The variability among these samples reflects the uncertainty about the underlying population.

The final row, labeled “real, bootstrap,” in Figure 3 shows the results of our Bayesian bootstrapping control procedure. We note that over 97% of our findings are reproduced in 100% of our Bayesian bootstrap iterations. For the remaining inconsistent three findings over the bootstrap, it is unfair to expect the private synthesizers to have higher epistemic parity than the bootstrap control.

Epistemic parity: overall performance. The overall performance of the synthesizers was impressive. All six synthesizers achieved 100% parity for Lee and Simpkins [38], and Fruith and Chan [24].

Figure 3: Epistemic parity for six competitive mechanisms for synthesizing data across four ϵ values ($e^{-3}, e^{-2}, e^{-1}, e^0, e^1, e^2$). All mechanisms achieve perfect parity on Fruiht and Chan [24] and Lee and Simpkins [38], and all but one achieve perfect parity on Pierce and Quiroz [46]. Only PATECTGAN can scale to support Jeong et al. [35]. All methods struggled with the high dimensionality of Iverson and Terry [31]. PrivMRF was too slow to be viable; we report results only for $\epsilon = e^0$. Only PrivBayes and MST achieved reasonable parity for Fairman et al. [23]. For datasets associated with Assari and Bazargan [2] and Saw et al. [59], only one finding was difficult to reproduce, and all methods struggled. Surprisingly, parity is relatively insensitive to ϵ .



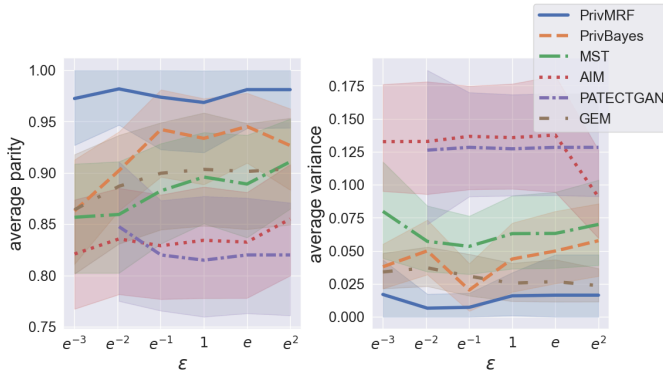
Besides PrivMRF (which was computationally infeasible to fit to the data), AIM, MST, PrivBayes, PATECTGAN, and GEM achieved 100% parity for Pierce and Quiroz [46] as well. Both Saw et al. [59], and Assari and Bazargan [2] also had very high levels of parity between findings on real and on synthetic data, although each of these papers had at least one finding that was difficult to reproduce.

Two of the papers provided the greatest challenge, and the most interesting results, across privacy regiments and synthesizer types: Fairman et al. [23], and Iverson and Terry [31]. These papers were challenging for very different reasons. Fairman et al. [23] had the second-smallest domain size, and the fewest variables. However, it had by far the largest sample, consisting of nearly 300K records. This combination made it very sensitive to noise in marginal measurements (as they are essentially counts), in turn making the findings difficult to replicate in low-privacy settings. Still, PrivBayes and

MST exhibited impressive performance in comparison to AIM, PATECTGAN and GEM. On the other hand, Iverson and Terry [31], had both one of the largest domains and the most variables of the papers in our benchmark, as well as a low mutual information between variables. No synthesizer with the exception of GEM exhibited convincing parity performance on this paper.

GEM was the strongest performing synthesizer on one paper (Iverson and Terry [31]). For the other papers in our benchmark, neither PATECTGAN nor GEM were the strongest performing. However, these methods were the most computationally tractable on high-dimensional large-domain data, and were the only methods that were feasible to run on Jeong et al. [35], where they both achieved 100% parity. Interestingly, PrivBayes often outperformed MST on our benchmark. We believe that this can be explained by two factors: (1) MST is tailored to work on high-dimensional

Figure 4: Average epistemic parity across papers achieved by AIM, PrivMRF, MST, PrivBayes, PATECTGAN, and GEM as a function of the privacy parameter $\epsilon \in \{e^{-3}, e^{-2}, e^{-1}, e^0, e^1, e^2\}$. Parity, on the y -axis, is on $[0,1]$ and represents the fraction of reproduced findings over all experiments at each ϵ .



datasets such as NIST, where explicitly parameterizing a conditional structure (like PrivBayes does) is costly and unstable, while the datasets in our benchmark are relatively low-dimensional; and (2) the findings that comprise the epistemic parity metric are based on conclusions that often rely on conditional relationships, which PrivBayes represents explicitly, while MST does not.

PrivMRF was the slowest synthesizer to run, and required a GPU. This requirement limited our ability to fully assess the capabilities of PrivMRF, although we observe that it performed well on the datasets on which it was able to run successfully. PrivBayes was the second-slowest method to run, due to a known limitation in handling high-dimensional data. However, PrivBayes performed competitively on datasets on which it was able to run successfully.

Notably, no synthesizer succeeded across all papers, and, remarkably, some findings were *never* reproduced by any of the synthesizers. We will revisit these findings further in Section 7.2.

Epistemic parity across ϵ values. Figure 4 compares synthesizer performance across reasonable ϵ values, shown on the x -axis in both sub-figures. The left side of the figure shows aggregated epistemic parity as the percentage of reproduced findings on the y -axis, over all iterations of each synthesizer, averaged over all publications in our benchmark. We observe that synthesizer performance (average parity) improves — although not substantially — for higher ϵ values for marginals-based methods PrivMRF, MST, and AIM. At the smallest values ($\epsilon = e^{-3}, e^{-2}$), the performance of PrivBayes, AIM, and MST all begin to noticeably (and understandably!) degrade, especially on certain findings (e.g., 16-21). Interestingly, PrivBayes achieves best performance at $\epsilon = e$, and PATECTGAN and GEM appear insensitive to the value of ϵ . These trends are consistent with the observations in Figure 3, and support the choice of $\epsilon = e$ as a reasonable privacy budget. Overall, we observe that restricting the privacy budget to $\epsilon = e^{-3}$ does not significantly affect the ability of the synthesizers to reproduce the “easy” findings, while increasing it to $\epsilon = e^2$ does not help with reproducing the “difficult” findings. We conjecture that the modeling structure employed by the synthesizer is more important than the scale of private noise.

Table 2: Methods used in benchmark papers, each corresponding to a type of *finding* in our framework.

	Descriptive Statistics	8
Regression	Between-Coefficients	4
	Fixed Coefficient (Sign)	2
	Variability	1
Causal Paths	Interaction	1
	Coefficient Difference	19
	PBR, FNR, FPR	2 (each)
Logistic Regression	Accuracy	2
	Mean Difference	24
	Temporal (FC)	26
Correlation	Pearson	12
	Spearman	1

The right side of Figure 4 shows average variance of epistemic parity. We observe that variance is lowest for PrivMRF, followed by PrivBayes. Further, we observe that the value of ϵ has little impact on parity variance; AIM is the only synthesizer that benefits from a higher value of ϵ in terms of reduced average parity variance.

The observation that epistemic parity is insensitive to ϵ is significant. It suggests that our metric is substantially different compared to other metrics that were previously used for assessment of DP synthesizers. Parity may provide insight into a more fundamental question about whether a DP synthesizer’s *methodology* — the types of measurements it takes to constitute a synthetic distribution — is appropriate to preserve the statistical properties of the dataset that are necessary to reproduce *findings*.

7.2 Qualitative results

Epistemic parity across finding types. Table 2 summarizes the methods used in the publications in our benchmark, each corresponding to a type of finding. We observe *Mean Difference* (both *Between-class* and *Temporal*) is by far the most common finding type, followed by *Coefficient Difference*. Whether a finding can be reproduced over DP synthetic data depends on several factors, including dataset size (as in Fairman et al. [23]) and dimensionality (as in Iverson and Terry [31]). However, finding type likely also plays a role: The majority (19 out of 26) of *Mean Difference / Temporal* findings are in these two papers that were difficult to reproduce. However, we must be cautious to interpret this as a trend: the remaining 7 findings of type *Mean Difference / Temporal (FC)* were in Saw et al. [59], and they were reproduced successfully by all synthesizers. In what follows, we qualitatively evaluate the impact of finding type (and, possibly, of other properties of the finding) on its reproducibility over DP synthetic data.

That some findings are easier to reproduce than others is unsurprising. Though each synthesizer relies on a fundamentally different approach to replicating the joint distribution across all of the data, they each struggle with high dimensional data. Further, for general-purpose synthetic data, PrivMRF, AIM, MST and PrivBayes prioritize lower dimensional 2- or 3-way relationships among variables, and thus it is unsurprising that simple mean comparison findings are easily preserved by these methods.

We were surprised at the high number of findings across all our papers (even those that we were unable to replicate) relying only on 1- or 2-dimensional comparisons: The low-dimensionality suggests

that earlier empirical studies (including Tao et al. [62] and Hay et al. [28]) may be suitable as proxy tasks. Targeted improvements to the synthesizers may allow us to simultaneously support high utility for individual findings and their composition into broad conclusions.

Next, we consider three findings that were difficult regardless of synthesizer or privacy regimen: #4 (Assari and Bazargan [2]), #39 (Iverson and Terry [31]), and #96 (Saw et al. [59]), see Figure 3.

Finding #4 is of type *Descriptive Statistics*. It is based on the text statement “Similarly, overall, people had 12.53 years of schooling at baseline (95%CI = 12.34-2.73).” Finding #39 is also of type *Descriptive Statistics*, and is based on a somewhat longer text statement that refers to specific percentages of individuals being diagnosed with specific disorders (5 such pairs of statistics in total). Finding #94 is of type *Mean Difference / Between-class*. It’s based on the text statement “From a longitudinal perspective, students from the two lower SES groups—low-middle and low SES groups—had significantly fewer persisters (31.9% and 29.9%) and emergers (6.1% and 5.4%) than their high SES peers (45.1% and 9.0%, respectively).”

These findings were difficult to reproduce because they give specific measurements for variables with large domains. Larger domains require proportionally more DP noise to privatize, and so the learned distribution over each of these variables was too noisy to reproduce these findings within our specified tolerance.

Visual findings. Figure 1, described in the Introduction, highlights a visual finding from Fairman et al. [23]. The results are subjectively similar, though a number of relationships may change at the individual level. Just as authors include visual findings in their papers, we argue that a DP evaluation should include qualitative results.

7.3 Summary of experimental results

Overall, we were encouraged by the performance of current state-of-the-art synthesizers on our benchmark. DP synthetic data has become more widely used in the social sciences (for Census Data, etc.) and these findings suggest that, in certain contexts, scientists can use DP synthetic data to conduct their scientific inquiry. We caveat this point: *Certain contexts* means relatively low-dimensional tabular data. Our benchmark can be used to assess if those data characteristics hold for a particular dataset, and researchers can proceed with their private analysis with increased confidence.

However, large domain and high-dimensional settings are still a challenge for DP synthesizers: as the domain/number of variables grows, the ease of *fingerprinting* individuals in a dataset increases dramatically. Our findings suggest that existing synthesizers struggle to scale (PrivMRF, MST, AIM, PrivBayes), or are far from achieving reasonable utility (PATECTGAN, GEM). We suggest incorporating more principled methods of data preprocessing, like DP-binning, DP variable pruning, or other domain/variable count reduction techniques into DP synthesizers, so that successful marginals-based methods can be utilized for more complex data.

8 CONCLUSIONS AND FUTURE WORK

Summary of contributions. In this paper, we proposed *epistemic parity* as methodology for measuring the utility of DP synthetic data in support of scientific research. We assembled a benchmark of peer-reviewed papers that analyze one of four studies in the ICPSR social science repository. We then experimentally evaluated

epistemic parity achieved by state-of-the-art DP synthesizers over the papers in our benchmark. Overall, we found epistemic parity to be a compelling method for evaluating DP synthesizers. Further, we found that, of the six DP synthesizers we evaluated, no single synthesizer outperformed all others on all papers. Finally, some findings were never reproduced by any of the synthesizers.

Despite facing well-known reproducibility challenges during benchmark construction, we are confident that our results will lead to generalizable insights. We are continuing to expand the benchmark with additional studies and papers, and invite others to use our open-source extensible framework, SynRD, to contribute.

Future work: Characterizing false discoveries. Replicating published findings using synthetic versions of the original data can reveal some implications of DP for scientific research. However, our epistemic parity methodology does not assess the possibility of findings that *would have occurred* if the original research had been done on synthetic data, or, more generally, the replicability of any hypotheses for which tests were not published. This challenge has been referred to as the file-drawer problem [32, 56] or publication bias. It may be a more severe issue when the original dataset contains more variables that could be used to formulate potential hypotheses. In future work, epistemic parity could be extended to quantify the effect of DP noise in producing false discoveries by simulating data with both “real” and spurious relationships.

Future work: Rebalancing utility and privacy. Though DP was developed to provide formal guarantees of privacy with best-effort utility, many practitioners and data providers may want the inverse: strong guarantees of utility with quantifiable, flexible risk of privacy violations that can be managed with policy and accountability rather than mathematical guarantees. We see this benchmark as promoting a more holistic discussion of socio-technical-legal systems for managing privacy-utility trade-offs. Besides, DP synthesizers, once already trained, can generate arbitrarily large samples at low cost. This makes the power of statistical hypothesis tests another concern for scientific research. Epistemic parity could be extended to use for calculating sample sizes, for example by Monte Carlo estimation of the sample size required to achieve a desired power for a particular finding, or for all findings in a given study.

Future work: Improving reproducibility and replicability of scientific discovery. Replicating published findings using synthetic versions of the original data can reveal some implications of DP for scientific research. However, our methodology does not assess the possibility of findings that *would have occurred* if the original research had been done on synthetic data, or, more generally, the replicability of any hypotheses for which tests were not published. This has been referred to as the file-drawer problem [32, 56] or publication bias. In future work, epistemic parity could be extended to quantify the effect of DP noise in producing findings that would not have been identified from the original data.

ACKNOWLEDGMENTS

This research was supported in part by NSF Awards Nos. 1916505, 1922658, 1934405, and NSF Graduate Research Fellowship Grant No. DGE-2039655, as well as Cisco Award 70618863, the Bill & Melinda Gates Foundation, and the NYU Center for Responsible AI.

REFERENCES

- [1] Edesio Alcobaca, Felipe Siqueira, Adriano Rivolli, Luis P. F. Garcia, Jefferson T. Oliva, and André C. P. L. F. de Carvalho. 2020. MFE: Towards reproducible meta-feature extraction. *Journal of Machine Learning Research* 21, 111 (2020), 1–5. <http://jmlr.org/papers/v21/19-348.html>
- [2] Shervin Assari and Mohsen Bazargan. 2019. Baseline obesity increases 25-year risk of mortality due to cerebrovascular disease: role of race. *International Journal of Environmental Research and Public Health* 16, 19 (2019), 3705.
- [3] Sergul Aydore, William Brown, Michael Kearns, Krishnaram Kenthapadi, Luca Melis, Aaron Roth, and Ankit A Siva. 2021. Differentially private query release through adaptive projection. In *International Conference on Machine Learning*. PMLR, 457–467.
- [4] Eugene Bagdasaryan, Omid Poursaeed, and Vitaly Shmatikov. 2019. Differential privacy has disparate impact on model accuracy. *Advances in neural information processing systems* 32 (2019).
- [5] Monya Baker. 2016. 1,500 scientists lift the lid on reproducibility. *Nature* 533, 7604 (2016).
- [6] Lorena A Barba. 2018. Terminologies for reproducible research. *arXiv preprint arXiv:1802.03311* (2018).
- [7] Claire McKay Bowen and Fang Liu. 2020. Comparative study of differentially private data synthesis methods. *Statist. Sci.* 35, 2 (2020), 280–307.
- [8] Danah Boyd and Jayshree Sarathy. 2022. Differential Perspectives: Epistemic Disconnects Surrounding the US Census Bureau’s Use of Differential Privacy. *Harvard Data Science Review (Forthcoming)* (2022).
- [9] Mark Bun and Thomas Steinke. 2016. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography: 14th International Conference, TCC 2016-B, Beijing, China, October 31–November 3, 2016, Proceedings, Part I*. Springer, 635–658.
- [10] Kuntai Cai, Xiaoyu Lei, Jianxin Wei, and Xiaokui Xiao. 2021. Data synthesis via differentially private markov random fields. *Proceedings of the VLDB Endowment* 14, 11 (2021), 2190–2202.
- [11] Colin F Camerer, Anna Dreber, Felix Holzmeister, Teck-Hua Ho, Jürgen Huber, Magnus Johannesson, Michael Kirchler, Gideon Nave, Brian A Nosek, Thomas Pfeiffer, et al. 2018. Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015. *Nature Human Behaviour* 2, 9 (2018), 637–644.
- [12] Miranda Christ, Sarah Radway, and Steven M Bellovin. 2022. Differential Privacy and Swapping: Examining De-Identification’s Impact on Minority Representation and Privacy Preservation in the US Census. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 1564–1564.
- [13] K Bretonnel Cohen, Jingbo Xia, Pierre Zweigenbaum, Tiffany J Callahan, Orin Hargraves, Foster Goss, Nancy Ide, Aurélie Névéol, Cyril Grouin, and Lawrence E Hunter. 2018. Three dimensions of reproducibility in natural language processing. In *LREC... International Conference on Language Resources & Evaluation: [proceedings]*. *International Conference on Language Resources and Evaluation*, Vol. 2018. NIH Public Access, 156.
- [14] Ben Dalton, Steven J Ingels, and Laura Fritch. 2016. *High School Longitudinal Study of 2009 (HLS:09). 2013 Update and High School Transcript Study: A First Look at Fall 2009 Ninth-Graders in 2013. NCES 2015-037rev*. Technical Report ICPSR36423.v1. Inter-University Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/ICPSR36423.v1>.
- [15] Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. 2021. Retiring adult: New datasets for fair machine learning. *NeurIPS* (2021).
- [16] Jiahao Ding, Xinyue Zhang, Xiaohuan Li, Junyi Wang, Rong Yu, and Miao Pan. 2020. Differentially private and fair classification via calibrated functional mechanism. In *AAAI*, Vol. 34. 622–629.
- [17] Irit Dinur and Kobbi Nissim. 2003. Revealing information while preserving privacy. In *Proceedings of the Twenty-Second ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 9–12, 2003, San Diego, CA, USA*, Frank Neven, Catriel Beeri, and Tova Milo (Eds.). ACM, 202–210. <https://doi.org/10.1145/773153.773173>
- [18] Dheeru Dua and Casey Graff. 2017. UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml>
- [19] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*. Springer, 265–284.
- [20] Cynthia Dwork, Moni Naor, Omer Reingold, Guy N Rothblum, and Salil Vadhan. 2009. On the complexity of differentially private data release: efficient algorithms and hardness results. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*. 381–390.
- [21] Cynthia Dwork, Aaron Roth, et al. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science* 9, 3-4 (2014), 211–407.
- [22] Timothy M Errington, Alexandria Denis, Nicole Perfito, Elizabeth Iorns, and Brian A Nosek. 2021. Reproducibility in cancer biology: challenges for assessing replicability in preclinical cancer biology. *Elife* 10 (2021), e67995.
- [23] Brian J Fairman, C Debra Furr-Holden, and Renee M Johnson. 2019. When marijuana is used before cigarettes or alcohol: Demographic predictors and associations with heavy use, cannabis use disorder, and other drug-related outcomes. *Prevention Science* 20, 2 (2019), 225–233.
- [24] Veronica Fruht and Thomas Chan. 2018. Naturally Occurring Mentorship in a National Sample of First-Generation College Goers: A Promising Portal for Academic and Developmental Success. 61, 3-4 (2018), 386–397. <https://doi.org/10.1002/ajcp.12233>
- [25] Georgi Ganev, Bristena Oprisanu, and Emiliano De Cristofaro. 2021. Robin Hood and Matthew Effects—Differential Privacy Has Disparate Impact on Synthetic Data. *arXiv preprint arXiv:2109.11429* (2021).
- [26] Moritz Hardt, Katrina Ligett, and Frank McSherry. 2010. A simple and practical algorithm for differentially private data release. *arXiv preprint arXiv:1012.4763* (2010).
- [27] Harris, Kathleen Mullan and Udry, J. Richard. 2022. *National Longitudinal Study of Adolescent to Adult Health (Add Health), 1994-2018 [Public Use]*. Technical Report ICPSR21600.v25. Inter-university Consortium for Political and Social Research [distributor], Carolina Population Center, University of North Carolina-Chapel Hill [distributor]. <https://doi.org/10.3886/ICPSR21600.v25>.
- [28] Michael Hay, Ashwin Machanavajjhala, Gerome Miklau, Yan Chen, and Dan Zhang. 2016. Principled evaluation of differentially private algorithms using dpbench. In *Proceedings of the 2016 International Conference on Management of Data*. 139–154.
- [29] Raquel Hill. 2015. Evaluating the Utility of Differential Privacy: A Use Case Study of a Behavioral Science Dataset. In *Medical Data Privacy Handbook*. Springer, 59–82.
- [30] James S. House. 2018. *Americans’ Changing Lives: Waves I, II, III, IV, and V, 1986, 1989, 1994, 2002, and 2011*. Technical Report ICPSR04690.v9. Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/ICPSR04690.v9>.
- [31] Grant L Iverson and Douglas P Terry. 2021. High school football and risk for depression and suicidality in adulthood: findings from a national longitudinal study. *Frontiers in neurology* 12 (2021).
- [32] Satish Iyengar and Joel B Greenhouse. 1988. Selection models and the file drawer problem. *Statist. Sci.* (1988), 109–117.
- [33] Matthew Jagielski, Michael Kearns, Jieming Mao, Alina Oprea, Aaron Roth, Saeed Sharifi-Malvajerdi, and Jonathan Ullman. 2019. Differentially private fair learning. In *ICML*. 3000–3008.
- [34] Bargav Jayaraman and David Evans. 2019. Evaluating differentially private machine learning in practice. In *28th USENIX Security Symposium (USENIX Security 19)*. 1895–1912.
- [35] Haewon Jeong, Michael D. Wu, Nilanjana Dasgupta, Muriel Médard, and Flavio P. Calmon. 2021. Who Gets the Benefit of the Doubt? Racial Bias in Machine Learning Algorithms Applied to Secondary School Math Education. In *35th Conference on Neural Information Processing Systems (NeurIPS 2021) Workshop on Math AI for Education (MATHAI4ED)*.
- [36] Christopher T Kenny, Shiro Kuriwaki, Cory McCartan, Evan TR Rosenman, Tyler Simko, and Kosuke Imai. 2021. The use of differential privacy for census data and its impact on redistricting: The case of the 2020 US Census. *Science advances* 7, 41 (2021), eabk3283.
- [37] Ronny Kohavi and Barry Becker. 1996. *UCI Mushroom Data Set*. Technical Report. UCI Machine Learning Repository. <https://archive.ics.uci.edu/ml/datasets/adult>
- [38] Glona Lee and Sandra D Simpkins. 2021. Ability self-concepts and parental support may protect adolescents when they experience low support from their math teachers. *Journal of Adolescence* 88 (2021), 48–57.
- [39] Terrance Liu, Giuseppe Vietri, and Steven Z Wu. 2021. Iterative methods for private synthetic data: Unifying framework and new methods. *Advances in Neural Information Processing Systems* 34 (2021), 690–702.
- [40] Ryan McKenna, Gerome Miklau, Michael Hay, and Ashwin Machanavajjhala. 2018. Optimizing error of high-dimensional statistical queries under differential privacy. *arXiv preprint arXiv:1808.03537* (2018).
- [41] Ryan McKenna, Gerome Miklau, and Daniel Sheldon. 2021. Winning the NIST Contest: A scalable and general approach to differentially private synthetic data. *arXiv preprint arXiv:2108.04978* (2021).
- [42] Ryan McKenna, Brett Mullins, Daniel Sheldon, and Gerome Miklau. 2022. AIM: An Adaptive and Iterative Mechanism for Differentially Private Synthetic Data. *arXiv preprint arXiv:2201.12677* (2022).
- [43] Marcus R Munafò, Brian A Nosek, Dorothy VM Bishop, Katherine S Button, Christopher D Chambers, Nathalie Percie du Sert, Uri Simonsohn, Eric-Jan Wagenmakers, Jennifer J Ware, and John Ioannidis. 2017. A manifesto for reproducible science. *Nature human behaviour* 1, 1 (2017), 1–9.
- [44] National Academies of Sciences, Engineering, and Medicine and others. 2019. Reproducibility and replicability in science. (2019).
- [45] Brian A Nosek, Charles R Ebersole, Alexander C DeHaven, and David T Mellor. 2018. The preregistration revolution. *Proceedings of the National Academy of Sciences* 115, 11 (2018), 2600–2606.
- [46] Kayla D. R. Pierce and Christopher S. Quiroz. 2019. Who matters most? Social support, social strain, and emotions. *Journal of Social and Personal Relationships* 36, 10 (2019), 3273–3292.

- [47] Haoyue Ping, Julia Stoyanovich, and Bill Howe. 2017. DataSynthesizer: Privacy-Preserving Synthetic Datasets. In *Proceedings of the 29th International Conference on Scientific and Statistical Database Management, Chicago, IL, USA, June 27-29, 2017*. ACM, 42:1–42:5. <https://doi.org/10.1145/3085504.3091117>
- [48] Fábio Pinto, Carlos Soares, and João Mendes-Moreira. 2016. Towards Automatic Generation of Metafeatures. In *Advances in Knowledge Discovery and Data Mining*, James Bailey, Latifur Khan, Takashi Washio, Gill Dobbie, Joshua Zhexue Huang, and Ruili Wang (Eds.). Springer International Publishing, Cham, 215–226.
- [49] Kristopher J. Preacher and Andrew F. Hayes. 2008. Asymptotic and Resampling Strategies for Assessing and Comparing Indirect Effects in Multiple Mediator Models. 40, 3 (2008), 879–891. <https://doi.org/10.3758/BRM.40.3.879>
- [50] Trivelloro E Raghunathan, Jerome P Reiter, and Donald B Rubin. 2003. Multiple imputation for statistical disclosure limitation. *Journal of official statistics* 19, 1 (2003), 1.
- [51] Ossi Räisä, Joonas Jälkö, Samuel Kaski, and Antti Honkela. 2022. Noise-Aware Statistical Inference with Differentially Private Synthetic Data. *arXiv preprint arXiv:2205.14485* (2022).
- [52] Adriano Rivolli, Luis Paulo F. Garcia, Carlos Soares, Joaquin Vanschoren, and André Carlos Ponce de Leon Ferreira de Carvalho. 2018. Characterizing classification datasets: a study of meta-features for meta-learning. *arXiv: Learning* (2018).
- [53] Lucas Rosenblatt, Joshua Allen, and Julia Stoyanovich. 2022. Spending Privacy Budget Fairly and Wisely. *CoRR abs/2204.12903* (2022). <https://doi.org/10.48550/arXiv.2204.12903> arXiv:2204.12903
- [54] Lucas Rosenblatt, Anastasia Holovenko, Taras Rumezhak, Andrii Stadnik, Bernease Herman, Julia Stoyanovich, and Bill Howe. 2022. Epistemic Parity: Reproducibility as an Evaluation Metric for Differential Privacy. *CoRR abs/2208.12700* (2022). <https://doi.org/10.48550/arXiv.2208.12700> arXiv:2208.12700
- [55] Lucas Rosenblatt, Xiaoyan Liu, Samira Pouyanfar, Eduardo de Leon, Anuj Desai, and Joshua Allen. 2020. Differentially private synthetic data: Applied evaluations and enhancements. *arXiv preprint arXiv:2011.05537* (2020).
- [56] Robert Rosenthal. 1979. The file drawer problem and tolerance for null results. *Psychological bulletin* 86, 3 (1979), 638.
- [57] Donald B Rubin. 1981. The bayesian bootstrap. *The annals of statistics* (1981), 130–134.
- [58] Steven Ruggles, Catherine Fitch, Diana Magnuson, and Jonathan Schroeder. 2019. Differential privacy and census data: Implications for social and economic research. In *AEA papers and proceedings*, Vol. 109. 403–08.
- [59] Guan Saw, Chi-Ning Chang, and Hsun-Yu Chan. 2018. Cross-sectional and longitudinal disparities in STEM career aspirations at the intersection of gender, race/ethnicity, and socioeconomic status. *Educational Researcher* 47, 8 (2018), 525–531.
- [60] Jeff Schlimmer. 1987. *UCI Adult Data Set*. Technical Report. UCI Machine Learning Repository. <https://archive.ics.uci.edu/ml/datasets/mushroom>.
- [61] Shun Takagi, Tsubasa Takahashi, Yang Cao, and Masatoshi Yoshikawa. 2021. P3gm: Private high-dimensional data release via privacy preserving phased generative model. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. IEEE, 169–180.
- [62] Yuchao Tao, Ryan McKenna, Michael Hay, Ashwin Machanavajjhala, and Gerome Miklau. 2021. Benchmarking Differentially Private Synthetic Data Generation Algorithms. *arXiv preprint arXiv:2112.09238* (2021).
- [63] Reihaneh Torkzadehmahani, Peter Kairouz, and Benedict Paten. 2020. DP-CGAN: Differentially Private Synthetic Data and Label Generation. *CoRR abs/2001.09700* (2020). arXiv:2001.09700 <https://arxiv.org/abs/2001.09700>
- [64] United States Department of Health and Human Services. 2016. *National Survey on Drug Use and Health (NSDUH), 2014*. Technical Report ICPSR36361.v1. Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/ICPSR36361.v1>.
- [65] Giuseppe Vietri, Grace Tian, Mark Bun, Thomas Steinke, and Steven Wu. 2020. New oracle-efficient algorithms for private synthetic data release. In *ICML*, 9765–9774.
- [66] Lei Xu, Maria Skoularidou, Alfredo Cuesta-Infante, and Kalyan Veeramachaneni. 2019. Modeling tabular data using conditional gan. *Advances in Neural Information Processing Systems* 32 (2019).
- [67] Jun Zhang, Graham Cormode, Cecilia M Procopiuc, Divesh Srivastava, and Xiaokui Xiao. SIGMOD 2014. PrivBayes: Private Data Release via Bayesian Networks. (SIGMOD 2014).