Assessing Utility of Differential Privacy for RCTs

Soumya Mukherjee [2] Araktrika Mustafi [2] Aleksandra Slavković [2] Lars Vilhuber[1]

[1]Labor Dynamics Institute, ILR, Cornell University, NY, United States

[2]Penn State University, PA, United States

May 2023

RCTs and Privacy

randomized control trials (RCTs)

have become a powerful tool for assessing the impact of interventions and policies in many contexts. Researchers have published an increasing number of studies that rely on RCTs for at least part of the inference.

Experimenter

is interested in determining the **main effect** of one or more **treatment variables** using a **regression model with fixed effects**.

May have additional variables used for **stratification** or blocking.

Treatment units within a block (unique combination of the blocking variables) are assigned to the treatment level combinations using **simple random sampling with replacement**.

Transparency and Privacy

Transparency

in the social sciences has lead supplementary materials to be made **public** as “replication packages”

RCTs and Privacy

De-identification as principal tool

Typical guidance followed by researchers who conduct RCTs [Department of Health and Human Services, 2012, Kopper, Sautmann and Turitto, 2020, DIME, 2020] suggests de-identification

Formal privacy methods?

- No guidance for social science researchers that is ”easy”
- Effect of implementation on inference unknown

This is particularly concerning because many of these studies have data from respondents in low and middle income countries (LMIC).

Research questions

Research question

Can DP be applied, can inference survive?

Publish data, maintain inference (possibly at some cost), do so easily.

Why RCTs?

Straightforward methods

OLS, difference-in-difference methods, possibly even simple difference in means across treated and untreated populations.

Most RCTs are small

samples of the overall population, allowing us (potentially) to leverage privacy-amplifying methods [Balle, Barthe and Gaboardi, 2018].

Make it real

Re-analysis of published RCTs

We analyze several previously published studies with **complete available data.**

Assess impact and cost

- Can inference validity be preserved?
- If not immediately, what is the cost (sample size, loss of precision) of achieving that?

Release protected data for others

For transparency, release protected data as part of (required) replication packages.

Data structure

Dataframe $M = [T, X]$

Dataframe with n rows and $p + t + b + 1$ columns, where

- n : total number of treatment units, $i = 1, \ldots, n$.
- t : number of dummy variables required for representing all possible treatment combination assignments, columns of T
- b : number of blocking variables.
- p : number of covariates, columns of X

Assumption

Covariates are sensitive

For now, we assume that the p **covariates** are sensitive, but do not address the b blocking variables or the t treatment conditions.

Model of the analyst

$$y_i = \alpha + \sum_{k=1}^{b} \tau_k T_{k,i} + \sum_{l=1}^{p} \gamma_l X_{l,i} + \nu_i, \quad i = 1, \ldots, n \tag{1}$$

where T_k represent the dummy variables for the treatment level combinations and X_l represent the covariates/control variables associated with the n treatment units and $\nu_i \overset{i.i.d}{\sim} N(0, \sigma^2)$. $beta = \{\alpha, \tau_k, \gamma_l\}$ are the parameters of the model, with τ_k of primary interest.

Model with blocks

When stratification is used with a total of m block combinations and n_j treatment units are assigned to j-th block combination, the corresponding regression model is given by

$$\begin{aligned} y_{ij} &= \alpha + \sum_{k=1}^{b} \tau_k T_{k,i} + \sum_{l=1}^{p} \gamma_l X_{l,ij} + \nu_{ij} \\ i &= 1, \ldots, n_j, j =, 1, \ldots, m, \sum_{j}^{m} n_j = n \end{aligned} \tag{2}$$

Parameters of interest

Treatment effects on outcomes

In both the above models, the parameter(s) of interest to the experimenter are the fixed effects τ_k, $k = 1, \ldots, b$.

Statistical utility

Inference concerning the fixed effects τ_k is affected as little as possible by the data release mechanism used to sanitize the analysis data in order to protect privacy.

Data Release Mechanism

Synthetic data

Synthetic dataframe with N observations that satisfies ϵ-differential privacy (DP).

Assumptions

Safe to release true (or close to true) treatment effects τ_k

Alternatively, one could separately protect the parameters of interest, but raises issues.

No need to protect assignment variables T

In practice, we either re-implement the design randomization mechanism, or sample without additional noise in the empirical distribution of post-random allocations, conditional on synthetic covariates.

Algorithm

Step 1

Construct a multivariate histogram for X.

- m = number of bins required to construct the histogram.
- C_i count/frequency of the observations in X corresponding to the i-th bin, $i = 1, \ldots, m$.
- C vector of counts given by $C = (C_1, \ldots, C_m)$.

Step 2: Laplace noise

- Draw m i.i.d observations $Z_1, \ldots, Z_m$ from a Laplace distribution with location parameter/mean 0 and variance $8/\epsilon^2$
- Compute sanitized vector of counts $D = (D_1, \ldots, D_m)$ where $D_i = C_i + Z_i$, $i = 1, \ldots, m$.
- Renormalize counts to be positive, $\widetilde{D} = (\widetilde{D}_1, \ldots, \widetilde{D}_m)$ where $\widetilde{D}_i = \frac{D_i \mathbf{I}_{D_i > 0}}{\sum_{i=1}^m D_i \mathbf{I}_{D_i > 0}}$, $i = 1, \ldots, m$.

Step 3: Create synthetic X

- Draw N i.i.d p-dimensional vectors $\widetilde{X}_1, \ldots, \widetilde{X}_N$ using simple random sampling with replacement from the m bins of the constructed histogram, with $\widetilde{D}$ as the corresponding probabilities of each of the m bins.
- Sanitized covariate dataframe = $\widetilde{X}^{N \times p} = \left[\widetilde{X}_1^T \ldots \widetilde{X}_N^T\right]^T$.

Step 4: Create synthetic T

- Construct T using the experimental design
- Alternatively, construct a similar but unprotected histogram as for X, sample
- $\widetilde{T}^{N \times t} = \left[\widetilde{T}_1 \ldots \widetilde{T}_t\right]$

We now have **synthetic dataframe** $\widetilde{M} = [\widetilde{T}, \widetilde{X}]$.

Step 5: Compute private $\hat{\beta}$

- Compute $\hat{\beta} = \{\hat{\tau}_k, \hat{\gamma}_l\}$ and $\hat{\sigma}^2$ using private M:
  $Y = \hat{\beta} M + \nu$, $\hat{sigma}^2 = var(\nu)$

Step 6: Impute $\widetilde{Y}$

- Construct $\widetilde{Y} = (\widetilde{Y}_1, \ldots, \widetilde{Y}_N)$ using privately computed $\hat{\beta}$ and $\hat{\sigma}^2$ as:

$$\widetilde{Y}_i = \widetilde{M}\hat{\beta} + Z_i$$

where $Z_i \overset{i.i.d}{\sim} N(0, \hat{\sigma}^2)$, $i = 1, \ldots, N$.

Step 7: Release $\widetilde{D}$

Release $\widetilde{D} = [\widetilde{Y}, \widetilde{M}] = [\widetilde{Y}, \widetilde{T}, \widetilde{X}]$.

Step 8: Estimate parameters $\widetilde{\tau}$

We can now estimate the parameters of interest on the protected (publicly available) data.

Proofs

The proof of differential privacy guarantee is based on Proposition 1 in Dwork et al. [2006] along with the post-processing property of pure differential privacy, while the statistical optimality is based on Theorem 4.4 of Wasserman and Zhou [2008].

Workflow

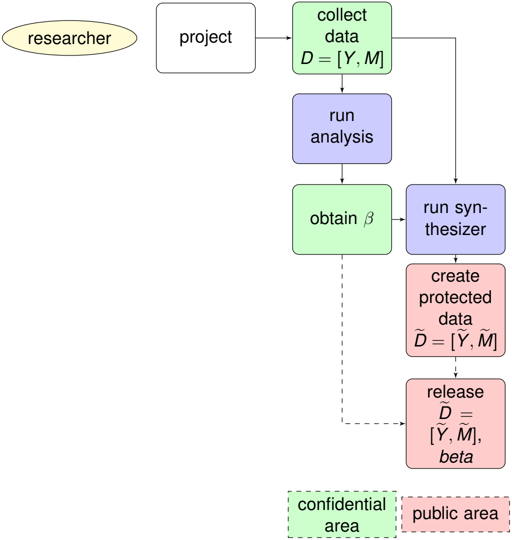
researcher
project
collect data $D = [Y, M]$
run analysis
obtain β
run synthesizer
create protected data $\widetilde{D} = [\widetilde{Y}, \widetilde{M}]$
release $\widetilde{D} = [\widetilde{Y}, \widetilde{M}]$, *beta*
confidential area
public area

Assessment

Given an unsanitized dataset $D = [Y, M]$. and a sanitized version of the same dataset (synthetic dataset) $\widetilde{D} = [\widetilde{Y}, \widetilde{M}]$ obtained using our proposed algorithm for a given privacy budget ϵ:

Metric 1 - C.I. overlap indicator:

computes whether there is any overlap between the 95% confidence intervals (C.I.) for the regression coefficients (individual C.I.'s for each regression coefficient)

Assessment

Metric 2 - Estimate coverage by sanitized C.I. indicator:
computes whether point estimates for the regression coefficients β_k computed based on the $D = [Y, M]$ fall within the confidence intervals for the regression coefficients computed based on the sanitized dataset $\widetilde{D} = [\widetilde{Y}, \widetilde{M}]$.

Metric 3

Metric 3: interval overlap measure J_k [Karr et al., 2006]

Consider the overlap of **confidence intervals** for variable n

- (L, U) for β_n (from the confidential data)

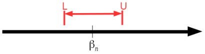

Metric 3

Metric 3: interval overlap measure J_k [Karr et al., 2006]

Consider the overlap of **confidence intervals** for variable n

- (L, U) for β_n (from the confidential data)
- $(\widetilde{L}, \widetilde{U})$ for $\widetilde{\beta}_n$ (from synthetic data)

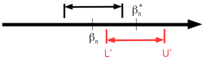

Metric 3

Metric 3: interval overlap measure J_k [Karr et al., 2006]

Consider the overlap of **confidence intervals** for variable n

- (L, U) for β_n (from the confidential data)
- $(\widetilde{L}, \widetilde{U})$ for $\widetilde{\beta}_n$ (from synthetic data)
- Let $L^{over} = \max(L, \widetilde{L})$ and $U^{over} = \min(U, \widetilde{U})$.

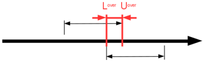

How well is inference preserved

Then the overlap in confidence intervals is

$$\widetilde{J}_k = \frac{1}{2}\left[\frac{U^{over} - L^{over}}{U - L} + \frac{U^{over} - L^{over}}{\widetilde{U} - \widetilde{L}}\right]$$

Metric 4

Metric 4 - Empirical Squared Error in Estimate:
$(\beta - \widetilde{\beta})^2$

Alternatively (not yet implemented)

Significant proximity of coefficients $t_{\Delta\beta_{k,m}}$

We compute

$$t_{\Delta\beta_{k,m}} = \frac{\beta_{k,m} - \widetilde{\beta}_{k,m}}{\sqrt{s_{k,m}^2 + \widetilde{s}_{k,m}^2}}$$

and assess its statistical significance. The fraction of insignificant tests across all estimated models and parameters is an indicator of how close the synthetic and confidential data are under the estimated models.

Assessing protection

In order to verify whether Aim 2 is satisfied, we choose a statistic that depends only on the sensitive data (the covariate data).

Metric 5 (Empirical Squared Error in Sensitive Statistic)

For a given statistic computed on M and $\widetilde{M}$, compute squared difference between the two values: $\Delta f = \left(f(M) - f(\widetilde{M})\right)^2$

Reported statistics

Average across simulations

Metrics 1 and 2 are reported as proportions, and Metrics 3, 4 and 5 are reported as averages, across multiple synthetic datasets/analyses. The goal is to obtain an indication of the performance of a single application of the proposed algorithm to obtain a single synthetic dataset, which is what we expect to be done in practice.

Baseline

Baseline variability

To distinguish variability introduced through the sampling and imputation from noise added by the protection mechanism (addition of Laplace noise), we perform the same synthetic data generation process, but without the addition of DP noise to the histogram counts (Step 2). This creates a "**synthetic unsanitized**" dataset $D^* = [Y^*, M^*]$. All metrics are calculated for D^* instead of $\widetilde{D}$ as well.

Numerical Experiments

Simulation Study 1

Simple dataframe

- $n = 100$ observations
- 1 treatment variable with two treatment levels, ”0” and ”1”, binomial distribution with equal probabilities
- $p = 1$ continuous covariate, Uniform(-5,5) or Beta(1,2).
- true regression coefficient as $\beta_0 = 0.05$ (Intercept term),$\beta_1 = 1, \beta_2 = 0.2$

Privacy parameters

Privacy budget

$\epsilon \in \{0.1, 0.5, 1\}$

Simulations

- For a given privacy budget, we simulate **n=100** different datasets D_n.
- For each D_n, we independently generate **s=20** synthetic datasets $\widetilde{D}_{n,\epsilon,s}$
- For each D_n, $\widetilde{D}_{n,\epsilon,s}$ we estimate the OLS model parameters and the five metrics.
- To compute Metric 5, we choose the variance of the covariate x_2

Simulations

Variable names	Metric 1	Metric 2	Metric 3	Metric 4
(Intercept)	1.00000	0.95000	0.79427	0.01127
x1	1.00000	0.94650	0.79718	0.02099
x2	1.00000	0.95350	0.78564	0.00076

Table: Effect on inference regarding regression coefficients measured using Metrics 1-4 for Simulation Study 1 with uniform covariate, averaged over 100 simulations of the sensitive dataframe, using 20 independently generated synthetic dataframes (with privacy budget $\epsilon = 0.1$) for each sensitive dataframe

Simulations

Variable names	Metric 1	Metric 2	Metric 3	Metric 4
(Intercept)	1.00000	0.95450	0.79703	0.01054
x1	1.00000	0.94600	0.79684	0.02094
x2	1.00000	0.94750	0.79166	0.00069

Table: Effect on inference regarding regression coefficients measured using Metrics 1-4 for Simulation Study 1 with uniform covariate, averaged over 100 simulations of the sensitive dataframe, using 20 independently generated synthetic dataframes (with privacy budget $\epsilon = 0.5$) for each sensitive dataframe

Simulations

Variable names	Metric 1	Metric 2	Metric 3	Metric 4
(Intercept)	1.00000	0.95000	0.79809	0.01046
x1	1.00000	0.94900	0.79737	0.02094
x2	1.00000	0.95700	0.79582	0.00065

Table: Effect on inference regarding regression coefficients measured using Metrics 1-4 for Simulation Study 1 with uniform covariate, averaged over 100 simulations of the sensitive dataframe, using 20 independently generated synthetic dataframes (with privacy budget $\epsilon = 1$) for each sensitive dataframe

Simulations

Privacy Budget	$\epsilon =$**0.1**	$\epsilon =$**0.5**	$\epsilon =$**1**	**Non-DP Synthesis**
MSE of Variance of x_2	6.888821	2.388792	1.273375	0.594822

Table: Effect on value of sensitive statistic (based on covariate data) measured using Metric 5 (MSE) for Simulation Study 1 using uniform covariate. Results are reported for DP synthesis with varying privacy budget ϵ and non-DP synthesis, each type of synthesis being averaged over 100 simulations of the sensitive dataframe, using 20 independently generated synthetic dataframes for each sensitive dataframe.

Application

Reducing Crime and Violence

Reducing Crime and Violence: Experimental Evidence from Cognitive Behavioral Therapy in Liberia

Christopher Blattman

Julian C. Jamison

Margaret Sheridan

AMERICAN ECONOMIC REVIEW
VOL. 107, NO. 4, APRIL 2017
(pp. 1165-1206)

Download Full Text PDF

Article Information

Abstract

We show that a number of noncognitive skills and preferences, including patience and identity, are malleable in adults, and that investments in them reduce crime and violence. We recruited criminally engaged men and randomized one-half to eight weeks of cognitive behavioral therapy designed to foster self-regulation, patience, and a noncriminal identity and lifestyle. We also randomized $200 grants. Cash alone and therapy alone initially reduced crime and violence, but effects dissipated over time. When cash followed therapy, crime and violence decreased dramatically for at least a year. We hypothesize that cash reinforced therapy's impacts by prolonging learning-by doing, lifestyle changes, and self-investment.

Reducing Crime and Violence: Table 2

Table 2—Program Impacts on Antisocial Behaviors

		ITT regression: ($N = 947$)											
		Therapy only				Cash only				Both			
				p-value				*p*-value				*p*-value	
Outcome	Control mean (1)	ITT (2)	SE (3)	Unadj. (4)	Adj. (5)	ITT (6)	SE (7)	Unadj. (8)	Adj. (9)	ITT (10)	SE (11)	Unadj. (12)	Adj. (13)
Panel A. 2–5 week impacts													
Antisocial behaviors, *z*-score	0.151	−0.249	[0.088]	**0.005**	**0.026**	−0.079	[0.091]	0.385	0.391	−0.308	[0.089]	**0.001**	**0.004**
Usually sells drugs	0.166	−0.077	[0.027]	**0.005**	0.082	−0.041	[0.029]	0.157	0.803	−0.076	[0.028]	**0.006**	0.099
# of thefts/robberies in past 2 weeks	2.577	−0.841	[0.400]	**0.036**	0.359	−0.770	[0.409]	0.060	0.502	−1.236	[0.407]	**0.002**	**0.045**
Disputes and fights in past 2 weeks, *z*-score	0.076	0.013	[0.092]	0.888	0.999	0.027	[0.091]	0.768	0.999	−0.132	[0.076]	0.084	0.596
Carries a weapon on body[a]	0.157	−0.086	[0.034]	**0.013**	0.167	−0.045	[0.037]	0.224	0.888	−0.093	[0.035]	**0.007**	0.115
Arrested in past 2 weeks	0.139	−0.011	[0.027]	0.674	0.999	0.006	[0.027]	0.819	0.999	−0.013	[0.029]	0.637	0.999
Aggressive behaviors, *z*-score	0.102	−0.208	[0.081]	**0.011**	0.151	0.008	[0.085]	0.928	0.999	−0.196	[0.087]	**0.024**	0.275
Verbal/physical abuse of partner, *z*-score[a]	−0.035	−0.087	[0.111]	0.430	0.985	0.091	[0.114]	0.422	0.985	−0.032	[0.115]	0.777	0.999
Panel B. 12–13 month impacts													
Antisocial behaviors, *z*-score	0.032	−0.083	[0.093]	0.373	0.878	0.132	[0.097]	0.175	0.681	−0.247	[0.088]	**0.005**	**0.037**
Usually sells drugs	0.135	−0.034	[0.029]	0.249	0.909	0.035	[0.030]	0.244	0.909	−0.059	[0.029]	**0.041**	0.474
# of thefts/robberies in past 2 weeks	1.839	0.073	[0.395]	0.855	0.997	0.352	[0.388]	0.365	0.943	−0.728	[0.363]	**0.045**	0.482
Disputes and fights in past 2 weeks, *z*-score	−0.060	−0.026	[0.091]	0.772	0.997	0.100	[0.090]	0.267	0.909	−0.100	[0.077]	0.194	0.881
Carries a weapon on	0.148	−0.059	[0.031]	0.061	0.553								

Reducing Crime and Violence: Table 2: Panel B

Panel B. 12–13 month impacts													
Antisocial behaviors, z-score	0.032	−0.083	[0.093]	0.373	0.878	0.132	[0.097]	0.175	0.681	−0.247	[0.088]	**0.005**	**0.037**
Usually sells drugs	0.135	−0.034	[0.029]	0.249	0.909	0.035	[0.030]	0.244	0.909	−0.059	[0.029]	**0.041**	0.474
# of thefts/robberies in past 2 weeks	1.839	0.073	[0.395]	0.855	0.997	0.352	[0.388]	0.365	0.943	−0.728	[0.363]	**0.045**	0.482
Disputes and fights in past 2 weeks, z-score	−0.060	−0.026	[0.091]	0.772	0.997	0.100	[0.090]	0.267	0.909	−0.100	[0.077]	0.194	0.881
Carries a weapon on body[a]	0.148	−0.059	[0.031]	0.061	0.553	0.043	[0.035]	0.215	0.894	−0.066	[0.033]	**0.049**	0.490
Arrested in past 2 weeks	0.118	−0.006	[0.024]	0.811	0.997	0.007	[0.025]	0.793	0.997	−0.033	[0.024]	0.171	0.863
Aggressive behaviors, z-score	0.188	−0.153	[0.110]	0.163	0.863	−0.043	[0.107]	0.685	0.997	−0.339	[0.109]	**0.002**	**0.035**
Verbal/physical abuse of partner, z-score[a]	−0.071	0.142	[0.100]	0.156	0.863	0.233	[0.113]	**0.040**	0.474	0.059	[0.104]	0.574	0.992

Reducing Crime and Violence: Table 2: Panel B: Main

This one:

Panel B. 12–13 month impacts													
Antisocial behaviors, z-score	0.032	−0.083	[0.093]	0.373	0.878	0.132	[0.097]	0.175	0.681	−0.247	[0.088]	**0.005**	**0.037**

In words

Specifically

long term effect of therapy and cash grant (12-13 months after the program) on a summary index of antisocial behaviours (referred to as `fam_asb_lt`) exhibited by a sample of 999 high-risk youths in Monrovia, Liberia.

Variables

Variable	Label
fam_asb_lt	ASB family index
cashassonly	Cash Only
tpassonly	Therapy Only
tpcashass	Both
tp_strata_alt	Therapy Block
cg_strata	Cash Block
age_b	Age
asbhostil_b	Barret ASB index
drugssellever_b	Drugs Sell indicator
drinkboozeself_b	Alcohol self indicator
druggrassself_b	Grass/Opium self indicator
harddrugsever_b	Hard Drugs indicator
steals_b	Steal self indicator

Effects, single dataset

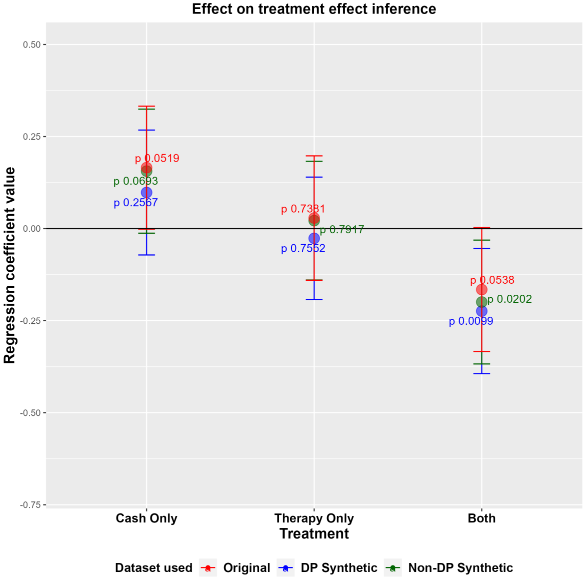

Effect numbers, single dataset

	M	$\widetilde{M}$	M^*
Cash Only	0.10	0.16	0.17
	(0.09)	(0.09)	(0.09)
Therapy Only	-0.03	0.02	0.03
	(0.08)	(0.08)	(0.09)
Both	-0.22^*	-0.20^*	-0.17
	(0.09)	(0.09)	(0.09)

$^*p < 0.05$

Average performance

Across 100 draws, with $\epsilon = 0.1, 0.5$ and 1,

same metrics as before. The variance of Age is taken as the sensitive statistic (which depends on the sensitive covariate data) for evaluating Metric 5.

Metrics Liberia study eps 0.1

Variable names	Metric 1	Metric 2	Metric 3	Metric 4
(Intercept)	1.00000	0.97000	0.80359	0.02076
Cash Only	1.00000	0.93000	0.79435	0.00816
Therapy Only	1.00000	0.92000	0.77931	0.00781
Both	1.00000	0.94000	0.80048	0.00720
Therapy Block	1.00000	0.98000	0.80903	0.00000
Cash Block	1.00000	0.95000	0.79263	0.00003
Age	1.00000	0.98000	0.73500	0.00001
Barret ASB index	1.00000	0.94000	0.74428	0.00027
Drugs Sell indicator	1.00000	0.95000	0.80788	0.00297
Alcohol self indicator	1.00000	0.96000	0.82439	0.00269
Grass/Opium self indicator	1.00000	0.97000	0.83612	0.00244
Hard Drugs indicator	1.00000	0.98000	0.81889	0.00271
Steal self indicator	1.00000	0.98000	0.82958	0.00254

Metrics Liberia study eps 0.5

Variable names	Metric 1	Metric 2	Metric 3	Metric 4
(Intercept)	1.00000	0.94000	0.80186	0.02202
Cash Only	1.00000	0.93000	0.79505	0.00817
Therapy Only	1.00000	0.92000	0.77917	0.00785
Both	1.00000	0.95000	0.79819	0.00732
Therapy Block	1.00000	0.98000	0.81052	0.00000
Cash Block	1.00000	0.96000	0.79342	0.00003
Age	1.00000	0.95000	0.73386	0.00001
Barret ASB index	1.00000	0.92000	0.74378	0.00028
Drugs Sell indicator	1.00000	0.94000	0.80342	0.00325
Alcohol self indicator	1.00000	0.95000	0.82324	0.00276
Grass/Opium self indicator	1.00000	0.97000	0.83083	0.00257
Hard Drugs indicator	1.00000	0.99000	0.82668	0.00235
Steal self indicator	1.00000	0.96000	0.81604	0.00293

Metrics Liberia study eps 1

Variable names	Metric 1	Metric 2	Metric 3	Metric 4
(Intercept)	1.00000	0.96000	0.80359	0.02092
Cash Only	1.00000	0.93000	0.79685	0.00801
Therapy Only	1.00000	0.93000	0.77940	0.00775
Both	1.00000	0.94000	0.79878	0.00722
Therapy Block	1.00000	0.98000	0.80874	0.00000
Cash Block	1.00000	0.96000	0.79245	0.00003
Age	1.00000	0.96000	0.73490	0.00001
Barret ASB index	1.00000	0.94000	0.74657	0.00027
Drugs Sell indicator	1.00000	0.95000	0.80468	0.00309
Alcohol self indicator	1.00000	0.96000	0.82834	0.00261
Grass/Opium self indicator	1.00000	0.98000	0.82588	0.00270
Hard Drugs indicator	1.00000	0.96000	0.81769	0.00264
Steal self indicator	1.00000	0.96000	0.81300	0.00306

Metrics Liberia study non DP

Variable names	Metric 1	Metric 2	Metric 3	Metric 4
(Intercept)	1.00000	0.95000	0.80074	0.03692
Cash Only	1.00000	0.98000	0.81815	0.00567
Therapy Only	1.00000	0.96000	0.83278	0.00478
Both	1.00000	0.94000	0.80123	0.00676
Therapy Block	1.00000	0.99000	0.82447	0.00000
Cash Block	1.00000	0.96000	0.77586	0.00003
Age	1.00000	0.94000	0.79168	0.00004
Barret ASB index	1.00000	0.93000	0.79746	0.00103
Drugs Sell indicator	1.00000	0.97000	0.80715	0.00618
Alcohol self indicator	1.00000	0.97000	0.79696	0.00447
Grass/Opium self indicator	1.00000	0.92000	0.79709	0.00470
Hard Drugs indicator	1.00000	0.97000	0.78692	0.00639
Steal self indicator	1.00000	0.95000	0.78054	0.00531

Privacy: Liberia study

Privacy Budget ϵ	**0.1**	**0.5**	**1**	**Non-DP**
MSE of Variance of Age	4481.74	4508.79	4503.16	0.9

Next steps

Next steps

Handle more variables.

We currently restricted ourselves to a few variables in a similar but not identical analysis to Blattman et al. Expand variables.

Next steps

Handle more variables.

We currently restricted ourselves to a few variables in a similar but not identical analysis to Blattman et al. Expand variables.

Handle multiple models

We want to release a single "plugin" protected dataset $[\widetilde{Y}, \widetilde{M}]$. But: $\widetilde{Y} = \hat{\beta}\widetilde{X}$ for a given model. Known tradeoffs - but how much in this context?

TABLE 2—PROGRAM IMPACTS ON ANTISOCIAL BEHAVIORS

		ITT regression: (N = 947)											
		Therapy only				Cash only				Both			
				p-value				p-value				p-value	
Outcome	Control mean (1)	ITT (2)	SE (3)	Unadj. (4)	Adj. (5)	ITT (6)	SE (7)	Unadj. (8)	Adj. (9)	ITT (10)	SE (11)	Unadj. (12)	Adj. (13)
Panel A. 2–5 week impacts													
Antisocial behaviors, z-score	0.151	−0.249	[0.088]	**0.005**	**0.026**	−0.079	[0.091]	0.385	0.391	−0.308	[0.089]	**0.001**	**0.004**
Usually sells drugs	0.166	−0.077	[0.027]	**0.005**	0.082	−0.041	[0.029]	0.157	0.803	−0.076	[0.028]	**0.006**	0.099
# of thefts/robberies in past 2 weeks	2.577	−0.841	[0.400]	**0.036**	0.359	−0.770	[0.409]	0.060	0.502	−1.236	[0.407]	**0.002**	**0.045**
Disputes and fights in past 2 weeks, z-score	0.076	0.013	[0.092]	0.888	0.999	0.027	[0.091]	0.768	0.999	−0.132	[0.076]	0.084	0.596
Carries a weapon on body[a]	0.157	−0.086	[0.034]	**0.013**	0.167	−0.045	[0.037]	0.224	0.888	−0.093	[0.035]	**0.007**	0.115
Arrested in past 2 weeks	0.139	−0.011	[0.027]	0.674	0.999	0.006	[0.027]	0.819	0.999	−0.013	[0.029]	0.637	0.999
Aggressive behaviors, z-score	0.102	−0.208	[0.081]	**0.011**	0.151	0.008	[0.085]	0.928	0.999	−0.196	[0.087]	**0.024**	0.275
Verbal/physical abuse of partner, z-score[a]	−0.035	−0.087	[0.111]	0.430	0.985	0.091	[0.114]	0.422	0.985	−0.032	[0.115]	0.777	0.999
Panel B. 12–13 month impacts													
Antisocial behaviors, z-score	0.032	−0.083	[0.093]	0.373	0.878	0.132	[0.097]	0.175	0.681	−0.247	[0.088]	**0.005**	**0.037**
Usually sells drugs	0.135	−0.034	[0.029]	0.249	0.909	0.035	[0.030]	0.244	0.909	−0.059	[0.029]	**0.041**	0.474
# of thefts/robberies in past 2 weeks	1.839	0.073	[0.395]	0.855	0.997	0.352	[0.388]	0.365	0.943	−0.728	[0.363]	**0.045**	0.482
Disputes and fights in past 2 weeks, z-score	−0.060	−0.026	[0.091]	0.772	0.997	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Carries a weapon on body[a]	[illegible]	[illegible]	[illegible]	[illegible]	0.553	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]

Next steps

Handle more variables.

We currently restricted ourselves to a few variables in a similar but not identical analysis to Blattman et al. Expand variables.

Handle multiple models

We want to release a single "plugin" protected dataset $[\widetilde{Y}, \widetilde{M}]$. But: $\widetilde{Y} = \hat{\beta}\widetilde{X}$ for a given model. Known tradeoffs - but how much in this context?

Expand to other papers

Our key interest is to show that this can be done (should be done) for arbitrary papers. Expanding to other RCTs is a first step.

Access and utility

Private data

needs to be access-limited. Not unique to this type of data (see IPUMS, PSID, etc.), but requires infrastructure

Re-use

is likely limited to re-analysis, but not broader re-use (may not be too limiting for RCTs)

Thank you!

Funding

NSF Award No. SES-1853209, CNS-1702760 & Bill & Melinda Gates Foundation, Subaward NO. 00011234 to The Pennsylvania State University

Bibliography

Abowd, John, and Ian M. Schmutte. 2015. “Economic analysis and statistical disclosure limitation.” Brookings Papers on Economic Activity, 221–267. https://doi.org/10.1353/eca.2016.0004.

Alabi, Daniel, Audra McMillan, Jayshree Sarathy, Adam Smith, and Salil Vadhan. 2020. “Differentially private simple linear regression.” https://arxiv.org/abs/2007.05157. tex.howpublished: arXiv:2007.05157 [cs.LG] tex.optabstract: tex.optgrants: Simons Investigator Award, Cooperative Agreement CB16ADR0160001 with the Census Bureau tex.optkeywords: tex.optsource:.

Awan, Jordan, and Aleksandra Slavković. 2020. “Structure and sensitivity in differential privacy: Comparing k-norm mechanisms.” Journal of the American Statistical Association, 1–20.

Balle, Borja, Gilles Barthe, and Marco Gaboardi. 2018. “Privacy amplification by subsampling: Tight analyses via couplings and divergences.” 6280–6290. http://papers.nips.cc/paper/7865-privacy-amplification-by-subsampling-tight-analyses-via-couplings-and-divergences. tex.bibsource: dblp computer science bibliography, https://dblp.org tex.biburl: https://dblp.org/rec/conf/nips/BalleBG18.bib tex.timestamp: Fri, 06 Mar 2020 17:00:31 +0100.

Barrientos, AndrÃ©s F., Alexander Bolton, Tom Balmat, Jerome P. Reiter, John M. de Figueiredo, Ashwin Machanavajjhala, Yan Chen, Charley Kneifel, and Mark DeLong. 2018. “Providing access to confidential research data through synthesis and verification: An application to data on employees of the U.S. federal government.” The Annals of Applied Statistics, 12(2): 1124 – 1156. https://doi.org/10.1214/18-AOAS1194.

Blattman, Christopher, Julian C. Jamison, and Margaret Sheridan. 2017. “Reducing Crime and Violence: Experimental Evidence from Cognitive Behavioral Therapy in Liberia.” American Economic Review, 107(4): 1165–1206. https://doi.org/10.1257/aer.20150503.

Blattman, Christopher, Julian C. Jamison, and Margaret Sheridan. n.d.. “Replication data for: Reducing Crime and Violence: Experimental Evidence from Cognitive Behavioral Therapy in Liberia.” https://doi.org/10.3886/E113056V1.

Bibliography II

Bowen, Claire McKay, Victoria Bryant, Leonard Burman, Surachai Khitatrakun, Robert McClelland, Philip Stallworth, Kyle Ueyama, and Aaron R Williams. 2020. “A synthetic supplemental public use file of low-income information return data: methodology, utility, and privacy implications.” 257–270, Springer.

Department of Health and Human Services. 2012. “Methods for De-identification of PHI.” `https://www.hhs.gov/hipaa/for-professionals/privacy/special-topics/de-identification/index.html` (accessed 2020-08-26).

DIME. 2020. “De-identification.” World Bank Dimewiki. `https://dimewiki.worldbank.org/De-identification` (accessed 2022-06-12).

Dwork, Cynthia, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. “Calibrating Noise to Sensitivity in Private Data Analysis.” Vol. Vol. 3876, 265–284. `https://doi.org/10.1007/11681878_14`.

Dwork, Cynthia, Frank McSherry, Kobbi Nissim, and Adam Smith. 2016. “Calibrating Noise to Sensitivity in Private Data Analysis.” Journal of Privacy and Confidentiality, 7(3). `https://doi.org/10.29012/jpc.v7i3.405`.

Dwork, Cynthia, Weijie Su, and Li Zhang. 2021. “Differentially private false discovery rate control.” Journal of Privacy and Confidentiality, 11(2). `https://doi.org/10.29012/jpc.755`.

Hundepool, Anco, Josep Domingo-Ferrer, Luisa Franconi, Sarah Giessing, Eric Schulte Nordholt, Keith Spicer, and Peter-Paul De Wolf. 2012. Statistical disclosure control. Vol. 2, Wiley New York.

Karr, A. F, C. N Kohnen, A Oganian, J. P Reiter, and A. P Sanil. 2006. “A Framework for Evaluating the Utility of Data Altered to Protect Confidentiality.” The American Statistician, 60(3): 224–232. `https://doi.org/10.1198/000313006X124640`.

Kopper, Sarah, Anja Sautmann, and James Turitto. 2020. “J-PAL GUIDE TO DE-IDENTIFYING DATA.” J-PAL. `https://www.povertyactionlab.org/sites/default/files/research-resources/J-PAL-guide-to-deidentifying-data.pdf` (accessed 2022-06-12).

Bibliography III

Machanavajjhala, Ashwin, Johannes Gehrke, Daniel Kifer, and Muthuramakrishnan Venkitasubramaniam. 2006. "l-Diversity: Privacy beyond k-Anonymity." 24. IEEE Computer Society. https://doi.org/10.1109/ICDE.2006.1. tex.bibsource: dblp computer science bibliography, https://dblp.org tex.biburl: https://dblp.org/rec/conf/icde/MachanavajjhalaGKV06.bib tex.timestamp: Wed, 16 Oct 2019 14:14:56 +0200.

Meager, Rachael. 2019. "Understanding the Average Impact of Microcredit Expansions: A Bayesian Hierarchical Analysis of Seven Randomized Experiments." American Economic Journal: Applied Economics, 11(1): 57–91. https://doi.org/10.1257/app.20170299.

Pistner, Michelle Nixon. 2020. Privacy Preserving Methods in the Era of Big Data: New Methods and Connections. https://etda.libraries.psu.edu/catalog/18340map5672.

Roth, Jonathan. 2022. "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends." American Economic Review: Insights, 4(3): 305–22. https://doi.org/10.1257/aeri.20210236.

Seeman, Jeremy, Aleksandra Slavkovic, and Matthew Reimherr. 2020. "Private Posterior Inference Consistent with Public Information: A Case Study in Small Area Estimation from Synthetic Census Data." 323–336, Springer.

Slavkovic, Aleksandra, and Jeremy Seeman. 2022. "Statistical Data Privacy: A Song of Privacy and Utility." https://doi.org/10.48550/ARXIV.2205.03336.

Slavkovic, Aleksandra, and Roberto Molinari. 2021. "Perturbed M-Estimation: A Further Investigation of Robust Statistics for Differential Privacy."

Vu, Duy, and Aleksandra Slavkovic. 2009. "Differential Privacy for Clinical Trial Data: Preliminary Evaluations." ICDMW '09, 138–143. Washington, DC, USA:IEEE Computer Society. https://doi.org/10.1109/ICDMW.2009.52.

Wasserman, Larry, and Shuheng Zhou. 2008. "A statistical framework for differential privacy." https://doi.org/10.48550/ARXIV.0811.2501.

Wood, Alexandra, Micah Altman, Kobbi Nissim, and Salil Vadhan. 2021. "Designing Access with Differential Privacy." In Handbook on Using Administrative Data for Research and Evidence-based Policy. , ed. Shawn Cole, Iqbal Dhaliwal, Anja Sautmann and Lars Vilhuber. Abdul Latif Jameel Poverty Action Lab. https://doi.org/10.31485/admindatahandbook.1.0.