Differential Privacy:

Observations for Economists

Daniel L. Goroff, Alfred P. Sloan Foundation

Not necessarily speaking for any of his institutional affiliations.

Some Public Goods Supported by Sloan

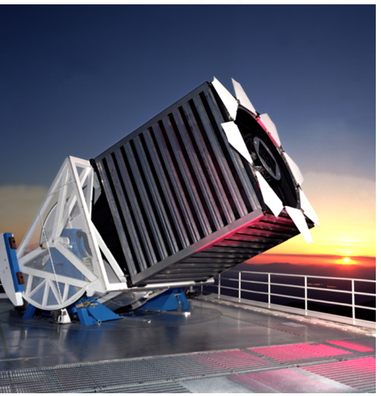

Economics: Paying Attention to Trade-Offs

E.g., Privacy vs. Accuracy. Here is a classification of protocols.

Obfuscation Stage	**Open Data**	**Data Enclave**	**Disclosure Agreement**	**Anonym-ization**	**Randomized Response**	**Multiparty Computing**	**Homomorph. Encryption**	**Differential Privacy**
Input			X		X	X	X	
Computation				X		X	X	X
Output		X	X				X	X

Privacy Protocol Promises?

No guarantees from these.

- **Data Enclave**: rely on enclave's discretion about what to release
- **Disclosure Agreement**: rely on provider's discretion about what to release
- **Anonymization**: rely on users discretion not try re-identification

Theorems say, "sanitizing data doesn't" and "de-identified data isn't."

(Dwork). See Dinur and Nissim (2003) on Database Reconstruction.

Randomized Response

Unbiased estimates of embarrassing information

- How many students in your class have ever cheated?
- Each flips a coin. Answer truthfully if heads.
- Otherwise flip again. Say yes if heads and no if tails.

Secure Multi-Party Computing

E.g. Find average of three peoples' salaries without revealing them

- Person i computes two random numbers, R_{ij} and R_{ik} .
- Then gives one to each of the other two people.
- Each takes salary, S_i ,subtracts his random numbers, adds those received.
- Result X_i can be shared. Their sum is the sum of salaries. Then divide by 3.

$$X_1 = S_1 - (R_{12} + R_{13}) + (R_{21} + R_{31})$$
$$X_2 = S_2 - (R_{21} + R_{23}) + (R_{12} + R_{32}) \qquad \Rightarrow X_1 + X_2 + X_3 = S_1 + S_2 + S_3$$
$$X_3 = S_3 - (R_{31} + R_{32}) + (R_{13} + R_{23})$$

Secure Multi-Party Computing

Challenges (Sloan support for such research since 2012)

- Other calculations may require bespoke algorithms.
- Passive attack is when parties collude but still follow the protocol.
- Active attack is where parties do not necessarily follow protocols.
- Can protect against passive if more than half are honest.
- Can protect against active attack if more than 2/3 are honest.
- Protection and functional evaluation may be computationally intensive.
- Releases the exact answer.

Fully Homomorphic Encryption

Computing on encrypted data (Sloan support since 2014)

- Suppose want to compute a function f of some sensitive data m.
- Encrypt your data using your own key to get enc(m).
- Without decrypting, can evaluate to get enc($f(m)$).
- Eval (Qx93aW, a2T5zN) = 78AbC3
- Multiparty, too, each with own key.
- Threshold decryption, e.g., to read result (exact).
- Very demanding computationally!

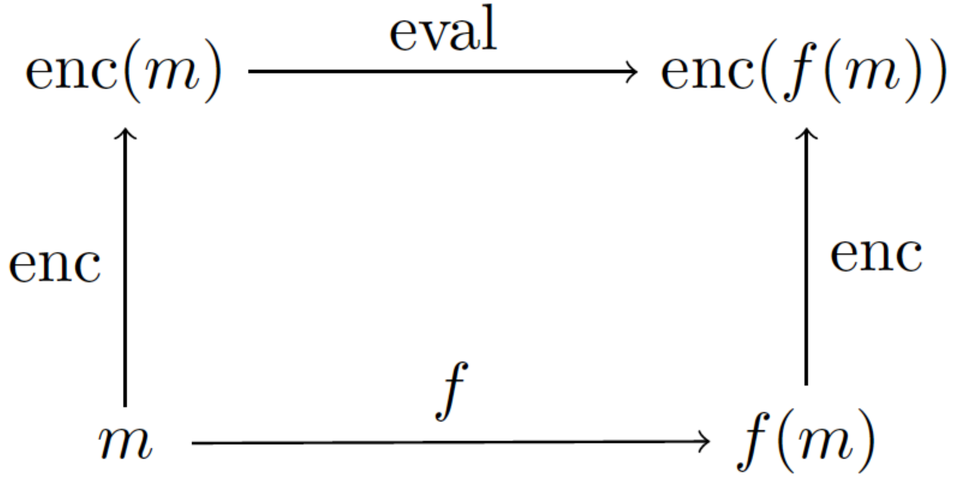

Differential Privacy

Motivation and Conceptual Framework

- A curator who answers too many questions too accurately will provably allow adversaries to reconstruct most of any database to a high degree of accuracy.
- To protect privacy, the curator's answer $M(x)$ as produced by a query mechanism M when applied to a dataset x, must be infused with noise.
- Want to limit how much an adversary can learn about whether the answer comes from a database $x = d$ that contains my information or from a "neighbor" $x = d'$ that is the same but missing the row with my information.
- The adversary has prior beliefs $\Pr(x = d)$ and $\Pr(x = d')$ about whether I am in the dataset, and updates those given $M(x)$ using Bayes' Law.

Economics: Bayesian Updating

How much can you learn about me by asking questions?

- Define "conditional probability" of event D given event A as $\Pr(D \mid A) = \Pr(D \cap A) / \Pr(A)$. It follows that:

$$\frac{\Pr(D \mid A)}{\Pr(D' \mid A)} = \frac{\Pr(A \mid D)}{\Pr(A \mid D')} \times \frac{\Pr(D)}{\Pr(D')}$$

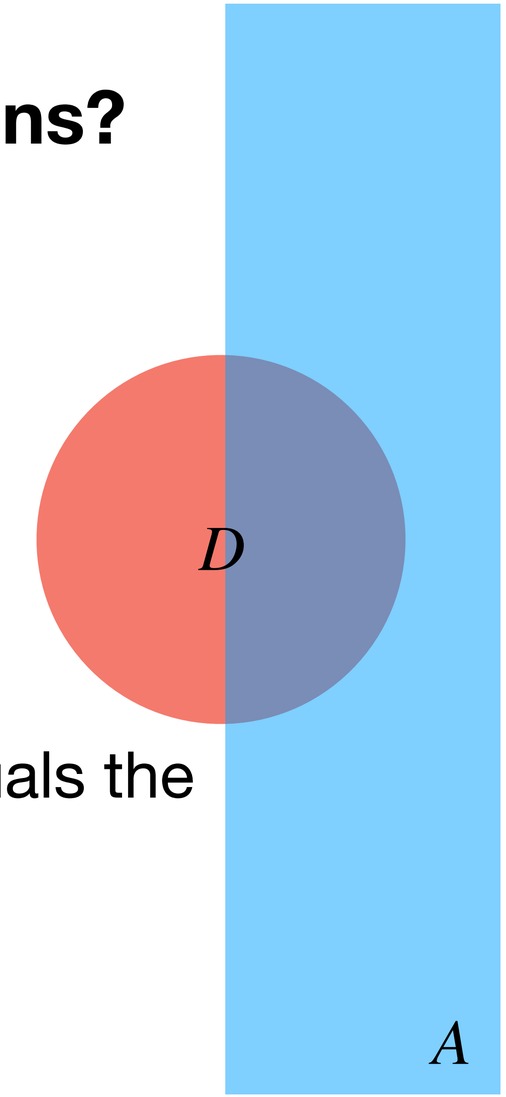

- Which says the posterior odds of D vs. D' once you know A equals the "Bayes Factor" times your prior odds of D vs. D'.

- So if the Bayes Factor is near 1, then A did not tell you much.

ϵ-Differential Privacy

When does a query mechanism satisfy this condition?

Let S be a subset of the image of a random mechanism M defined on datasets. We want to limit how much your prior odds about whether $x = d$ or d' can change by learning that $M(x)$ belongs to S. Bayes Law says:

$$\frac{\Pr(x = d \,|\, M(x) \in S)}{\Pr(x = d' \,|\, M(x) \in S)} = \frac{\Pr(M(x) \in S) \,|\, x = d)}{\Pr(M(x) \in S) \,|\, x = d')} \times \frac{\Pr(x = d)}{\Pr(x = d')} \quad \text{or}$$

$$\frac{\Pr(x = d \,|\, M(x) \in S)}{\Pr(x = d' \,|\, M(x) \in S)} = \frac{\Pr(M(d) \in S))}{\Pr(M(d') \in S))} \times \frac{\Pr(x = d)}{\Pr(x = d')}$$

So we keep that Bayes Factor near 1 by requiring for neighboring datasets d and d' :

$$\frac{\Pr(M(d) \in S))}{\Pr(M(d') \in S))} \leq \exp(\epsilon) \approx 1 + \epsilon + \ldots$$

Do Such Algorithms Exist?

Laplace Mechanism is one example.

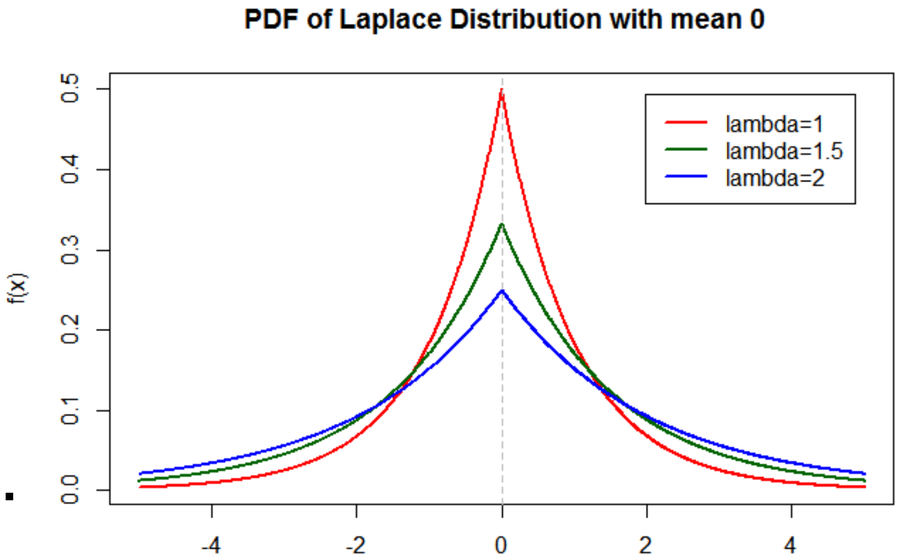

- Let f be a deterministic function defined on datasets.
- Let Y denote a Laplacian random variable whose density at t is $\propto \exp(|t|/\lambda)$.
- Then the mechanism $M(x) = f(x) + Y$ satisfies ϵ-DP for $\lambda = \Delta f/\epsilon$, where:
- $\Delta f = \max |f(d) - f(d')|$ over all neighboring d and d' denotes the sensitivity of f.
- For counting queries in particular, notice that $\Delta f = 1$. (Regression coefficients?)
- More $\epsilon \sim$ less privacy & more accuracy. Small $\epsilon \sim$ more privacy & less accuracy.

Properties Implied by DP

Suppose mechanism M_i satisfies ϵ_i - differential privacy

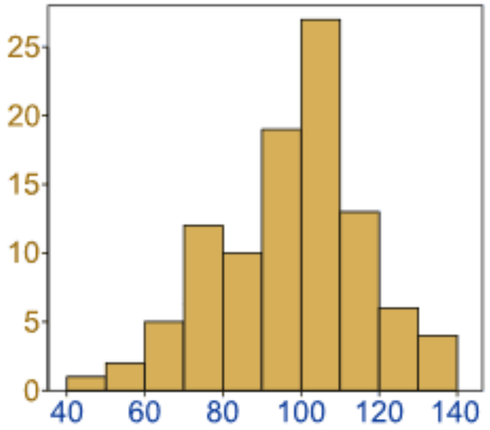

- Should I allow my information to be included in a study? If the research question is answered by mechanism M_1, then I am guaranteed no one's prior odds that my data is even there can change by more than a factor of $\exp(\epsilon_1) \approx 1 + \epsilon_1 + \ldots$
- That guarantee is immune to post-processing: i.e., if G is a random or deterministic function, $G(M_1)$ still satisfies ϵ_1-differential privacy.
- Sequential Composition: $g(M_1, M_2)$ satisfies $(\epsilon_1 + \epsilon_2)$-differential privacy.
- Parallel Composition: If $M_1(x)$ and $M_2(x)$ are always computed on disjoint elements of x, then $g(M_1, M_2)$ satisfies $\max\{\epsilon_1, \epsilon_2\}$-differential privacy.

DP Features and Challenges

No other formal frameworks available (Sloan since 2011)

- DP provides a formal privacy guarantee. You can still be harmed based on a study's results, but it protects against harm due to participating in the study.
- DP not only provides a conceptual framework, it also has composition rules and a clear sense of the trade-offs along with an adjustment parameter, ϵ.
- Fixing ϵ is a policy question, not a mathematical one. Once set for x, if you want to answer more than one question, say M_1 and M_2, then you have to make sure $(\epsilon_1 + \epsilon_2) \leq \epsilon$. In other words, ϵ imposes a "privacy loss budget."
- Note that answering even an innocent looking question about x without first infusing noise blows the privacy loss budget and voids the privacy guarantees.

Economics: Evidence is a Rival Good

Research Inevitably Leaks Privacy and Validity

- So if you try cajoling an agency, business, or individual into submitting sensitive data—encrypted or otherwise—but still plan to release precise statistics, it is not clear why they should trust you to protect their privacy.
- *Every query answered leaks some privacy.* DP controls the rate, though.
- *Every query answered leaks validity, too*, by getting closer to “p-hacking." Again, DP controls the rate of overfitting. Why sacrifice *any* privacy for junk science?
- So if you don’t like using an ϵ-DP mechanism for adaptive data analysis, then you don’t mind overfitting. And if you can’t handle noise *whose distribution you know*…
- Economically, empirical evidence from (private) data is like other commodities *that you use up*. This is a theorem, not something fixable by new technologies.

End-to-End Solutions

Like privacy protection in other realms.

- Will need MPC + FHE + DP to cover collection, processing, and output stages. Model selection for research, synthetic data for rummaging, reference data, too. Many are working on such integrated suites, including companies and countries.
- Sloan Bets: Actuate, a nonprofit founded for this purpose by a former DARPA director, and OpenDP, an open source development project based at Harvard.
- Especially if giving up some privacy is inevitable, it needs to be worth it. So success criteria include not just algorithm implementation but the capacity of the entire system to put data to good use while protecting privacy throughout.
- Administrative data, like that from agencies or businesses, can be especially difficult to use well in this way. One of many challenges is data linkage.

Privacy-Protecting Data Linkage

Including both entity resolution and de-duplication

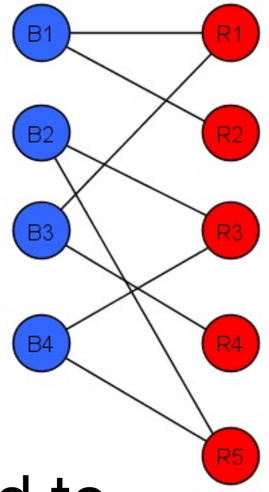

- Real data always contains errors, ambiguities, and missing fields. Need lots of other fields to try to match two records that are supposed to describe the same thing. Identifying information is naturally most useful.
- How can this be done while protecting privacy? Can machines do it for us?
- Old methods rely on testing the hypothesis that two records belong linked. Involves arbitrary thresholds. Results are not transitive. Don't scale well.
- Machine-Learning algorithms are not yet very good at this either.
- Sloan Bet: Bayesian optimization using MCMC over links in bipartite graphs connecting records to latent sources. See R. Steorts at Duke.

Administrative Data Challenges

Economics: Institutions can lower transaction costs

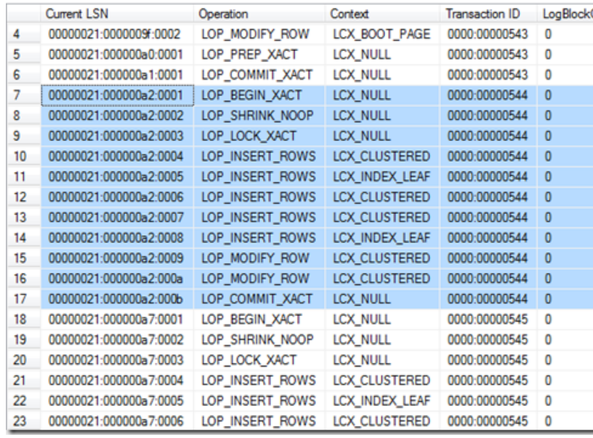

	Current LSN	Operation	Context	Transaction ID	LogBlock
4	00000021:0000009f:0002	LOP_MODIFY_ROW	LCX_BOOT_PAGE	0000:00000543	0
5	00000021:000000a0:0001	LOP_PREP_XACT	LCX_NULL	0000:00000543	0
6	00000021:000000a1:0001	LOP_COMMIT_XACT	LCX_NULL	0000:00000543	0
7	00000021:000000a2:0001	LOP_BEGIN_XACT	LCX_NULL	0000:00000544	0
8	00000021:000000a2:0002	LOP_SHRINK_NOOP	LCX_NULL	0000:00000544	0
9	00000021:000000a2:0003	LOP_LOCK_XACT	LCX_NULL	0000:00000544	0
10	00000021:000000a2:0004	LOP_INSERT_ROWS	LCX_CLUSTERED	0000:00000544	0
11	00000021:000000a2:0005	LOP_INSERT_ROWS	LCX_INDEX_LEAF	0000:00000544	0
12	00000021:000000a2:0006	LOP_INSERT_ROWS	LCX_CLUSTERED	0000:00000544	0
13	00000021:000000a2:0007	LOP_INSERT_ROWS	LCX_CLUSTERED	0000:00000544	0
14	00000021:000000a2:0008	LOP_INSERT_ROWS	LCX_INDEX_LEAF	0000:00000544	0
15	00000021:000000a2:0009	LOP_MODIFY_ROW	LCX_CLUSTERED	0000:00000544	0
16	00000021:000000a2:000a	LOP_MODIFY_ROW	LCX_CLUSTERED	0000:00000544	0
17	00000021:000000a2:000b	LOP_COMMIT_XACT	LCX_NULL	0000:00000544	0
18	00000021:000000a7:0001	LOP_BEGIN_XACT	LCX_NULL	0000:00000545	0
19	00000021:000000a7:0002	LOP_SHRINK_NOOP	LCX_NULL	0000:00000545	0
20	00000021:000000a7:0003	LOP_LOCK_XACT	LCX_NULL	0000:00000545	0
21	00000021:000000a7:0004	LOP_INSERT_ROWS	LCX_CLUSTERED	0000:00000545	0
22	00000021:000000a7:0005	LOP_INSERT_ROWS	LCX_INDEX_LEAF	0000:00000545	0
23	00000021:000000a7:0006	LOP_INSERT_ROWS	LCX_CLUSTERED	0000:00000545	0

- "Administrative Data" refers to information not originally collected for research purposes, e.g., the transaction records of agencies, businesses, or people.
- Besides privacy and linking challenges, many other factors make such data very difficult to use well. Also need to deal privately with cleaning, metadata, model selection, missing data, selection bias, and very high transactions costs.
- ADP only makes it look easy to predict the BLS unemployment rate figures. Risks include capture, front running, entry, exit, accounting changes, manipulation, etc.
- When transaction costs are high, economics says the solutions are institutional. Sloan bet: Administrative Data Research Institute as a network of sector-specific data intermediaries run by A. O'Hara at Georgetown.

National Secure Data Service

Called for by the bipartisan Evidence Act of 2018

- Set up as an FFRDC run by a research agency like NSF, such an NSDS could share data with companies and agencies under CIPSEA protection.
- So behind its firewall, NSDS could assemble the “mother of all statistical frames.” Census data already underlies nearly all such sampling frames. This is as important a use case as nearly anything else the Census does.
- Such a frame could make linking much simpler, and also provide vital weights to adjust for the selection bias in non-representative survey or administrative data.
- Use MPC + FHE + DP for collection, processing, and release protocols.
- Sloan Bet: The Data Foundation and COPAFS are conducting design exercises.

Observations about DP

Basic concepts are familiar to economists:

- Trade-Offs
- Bayes Factors
- Modeling with Noisy Data
- Allocation of Scare Resources

Now that we know why both privacy and validity are such precious commodities, we need even more ideas to help facilitate empirical economics research accordingly.

References

1. Goroff, D.L., 2015. Balancing privacy versus accuracy in research protocols. *Science*, *347*(6221), pp.479-480.
2. Dwork, C. and Roth, A., 2014. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, *9*(3–4), pp.211-407.
3. Gaboardi, M., Hay, M. and Vadhan, S., 2020. A Programming Framework for OpenDP.
4. Steorts, R.C., Hall, R. and Fienberg, S.E., 2016. A Bayesian approach to graphical record linkage and deduplication. *Journal of the American Statistical Association*, *111*(516), pp.1660-1672.
5. Bailey, M., Cole, C., Henderson, M. and Massey, C., 2018. How Well Do Automated Linking Methods Perform in Historical Data? Evidence from New US Ground Truth.
6. Dwork, C., Feldman, V., Hardt, M. Pitassi, T., Reingold, O. and Roth, A.L., 2015, June. Perserving statistical validity in adaptive data analysis. In *Proceedings of the 47th annual ACM symposium on theory of computing* (pp. 117-126)
7. Goroff, D., Polonetsky, J. and Tene, O., 2018. Privacy protective research: Facilitating ethically responsible access to administrative data. *The ANNALS of the American Academy of Political and Social Science*, *675*(1), pp.46-66.

Differential Privacy:

Observations for Economists

Daniel L. Goroff, Alfred P. Sloan Foundation

Not necessarily speaking for any of his institutional affiliations.