What Is a Privacy-Loss Budget and How Is It Used to Design Privacy Protection for a Confidential Database?

John M. Abowd
Associate Director for Research and Methodology and Chief Scientist
U.S. Census Bureau

ASA Privacy Day Webinar
February 1, 2018

United States Census Bureau | U.S. Department of Commerce
Economics and Statistics Administration
U.S. CENSUS BUREAU
census.gov

Acknowledgments and Disclaimer

- This webinar is based on collaborative work with Ian Schmutte, University of Georgia [link], and Lars Vilhuber
- The application to Census Bureau publication incorporates work by Dan Kifer (Scientific Lead), Simson Garfinkel (Senior Scientist for Confidentiality and Data Access), Tammy Adams, Robert Ashmead, Michael Bentley, Stephen Clark, Aref Dajani, Jason Devine, Michael Hay, Cynthia Hollingsworth, Michael Ikeda, Philip Leclerc, Ashwin Machanavajjhala, Gerome Miklau, Brett Moran, Edward Porter, and Anne Ross [link to the September 2018 Census Scientific Advisory Committee presentation]
- Parts of this talk were supported by the National Science Foundation, the Sloan Foundation, and the Census Bureau (before and after my appointment started)
- The opinions expressed in this talk are the my own and not necessarily those of the U.S. Census Bureau

Outline

- First, a little economics
- Then, a little computer science
- And, put them together
- Two examples
- Thinking about the choice problem for local differential privacy
- Thinking about the choice problem for global differential privacy
- Now, the challenge for the 2020 Census of Population and Housing

A Little Economics

- Finite resource: information in an existing database
- Competing uses:
  - Accuracy (fitness-for-use) of published statistics
  - Loss of privacy (information leakage about individuals)
- Optimal resource allocation should equate:
  - Marginal rate of transformation (opportunity cost)
  - Marginal willingness to pay (marginal rate of substitution)
- Both accuracy and privacy are public goods
- Private provision will generally produce sub-optimal accuracy or privacy loss

Production Possibilities

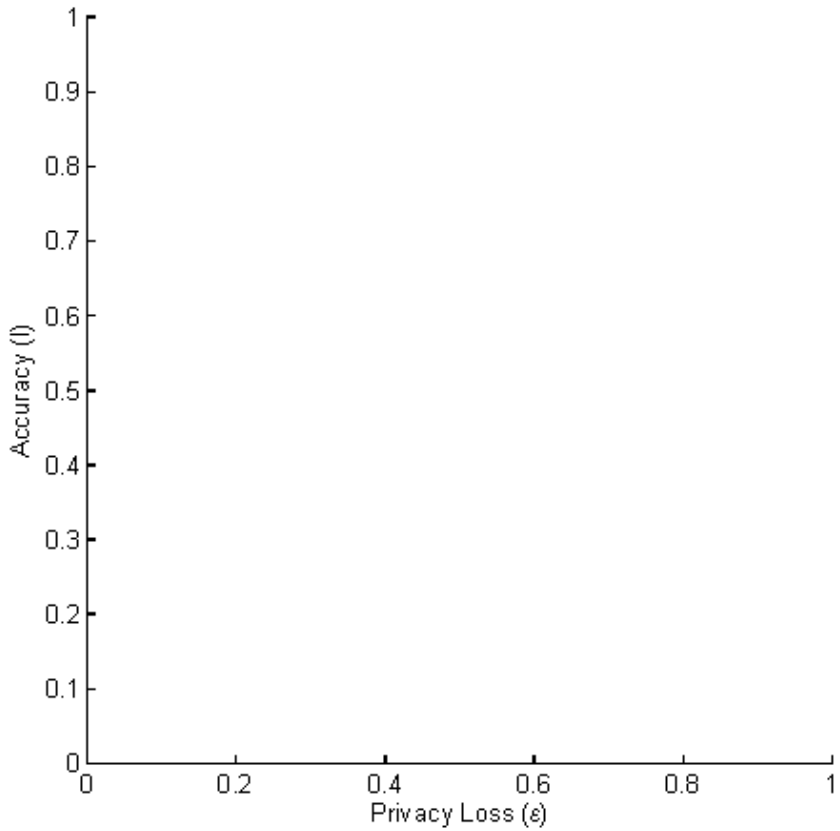

Production Possibilities

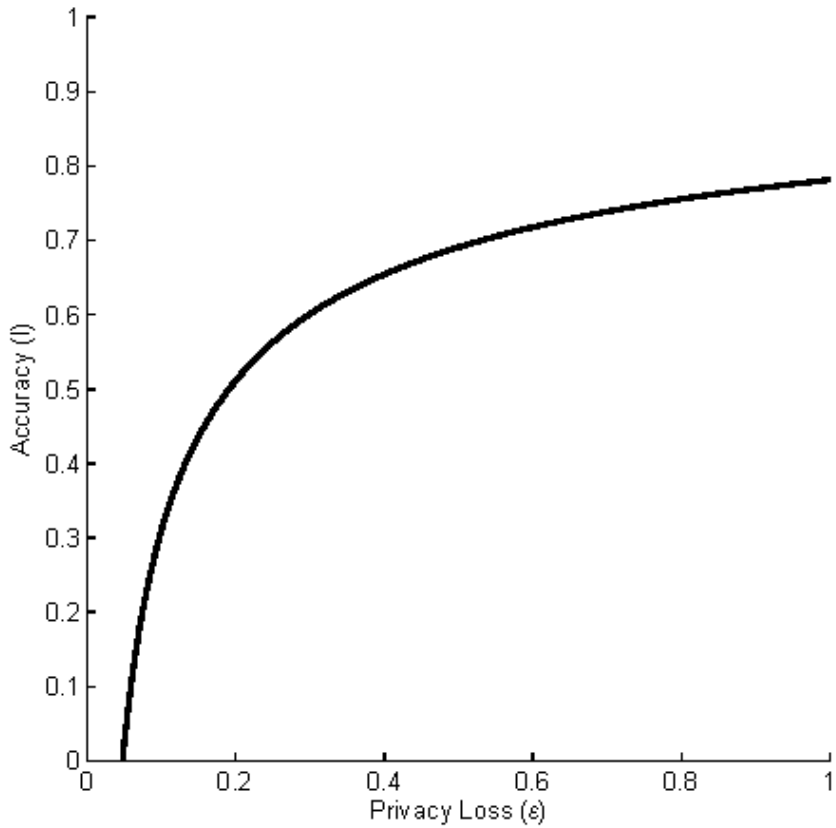

Production Possibilities

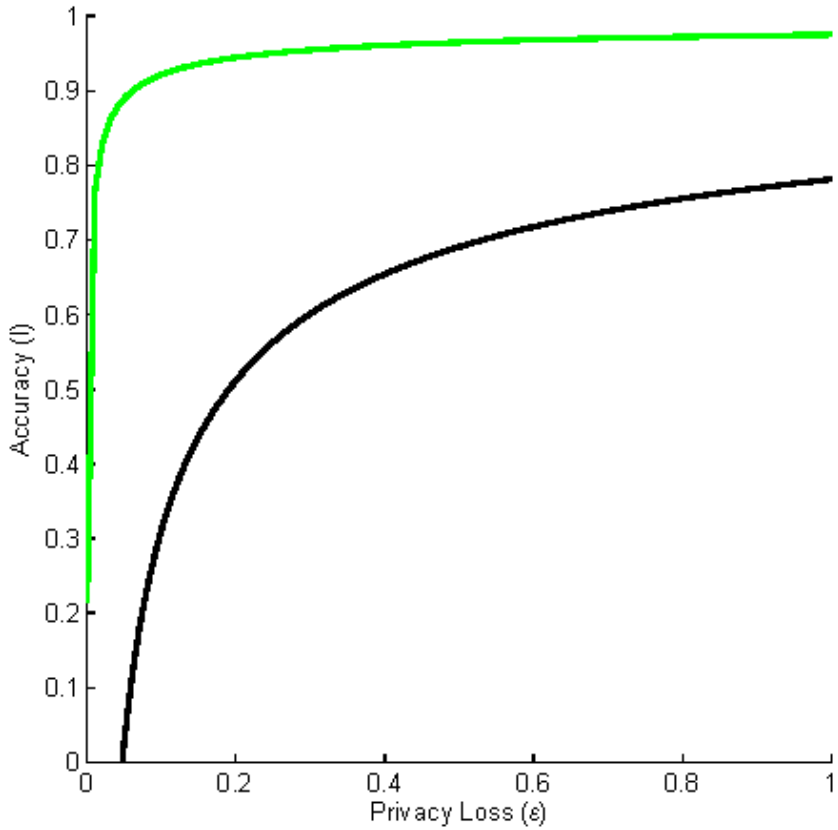

Social Welfare Function

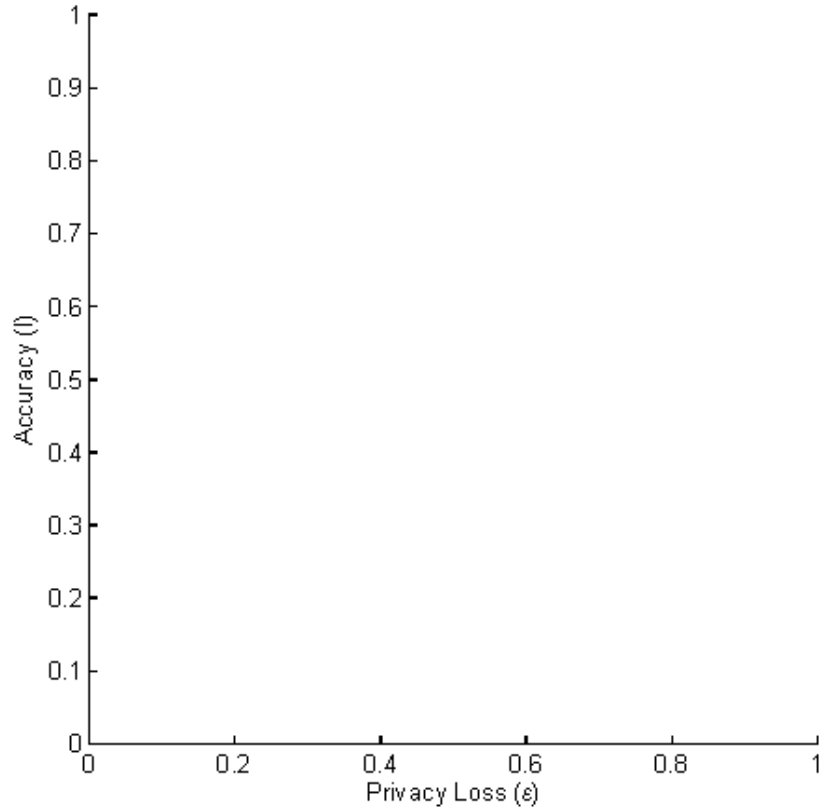

Social Welfare Function

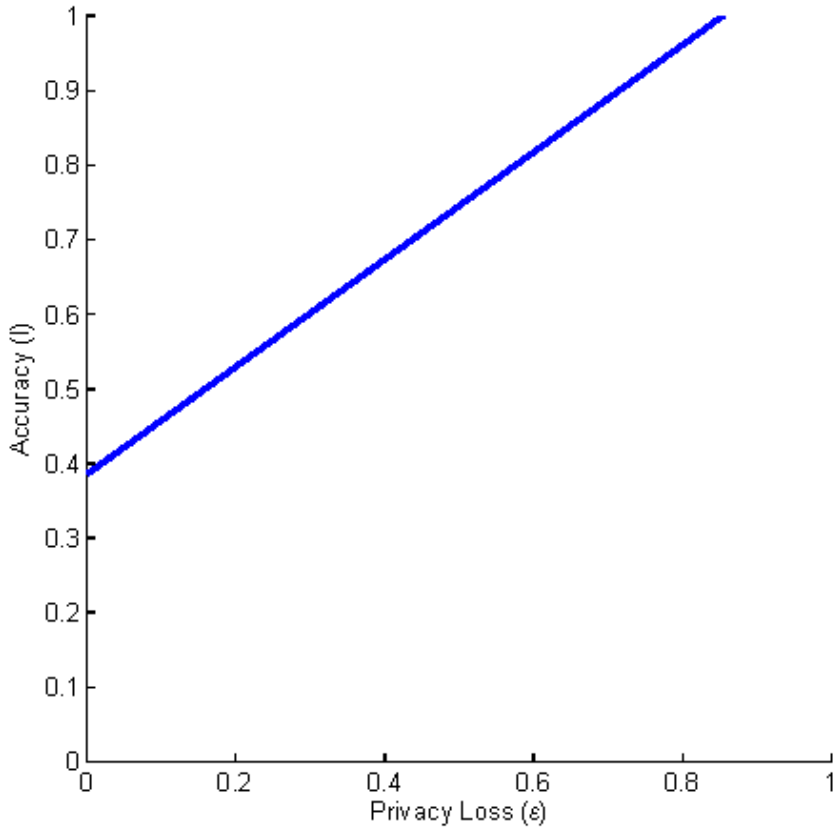

Social Welfare Function

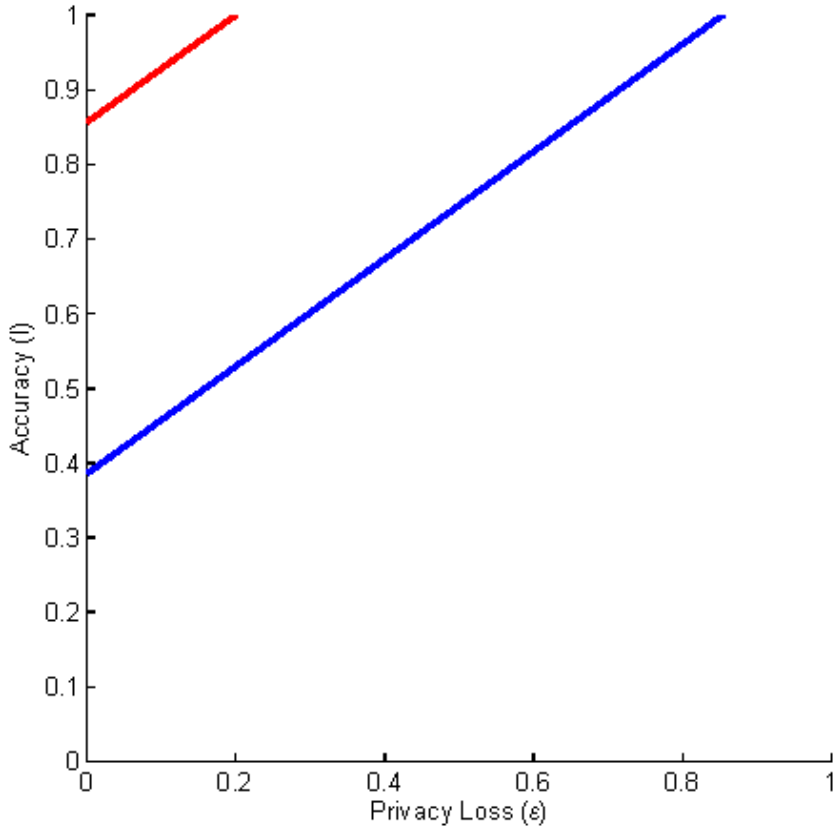

Social Welfare Maximization

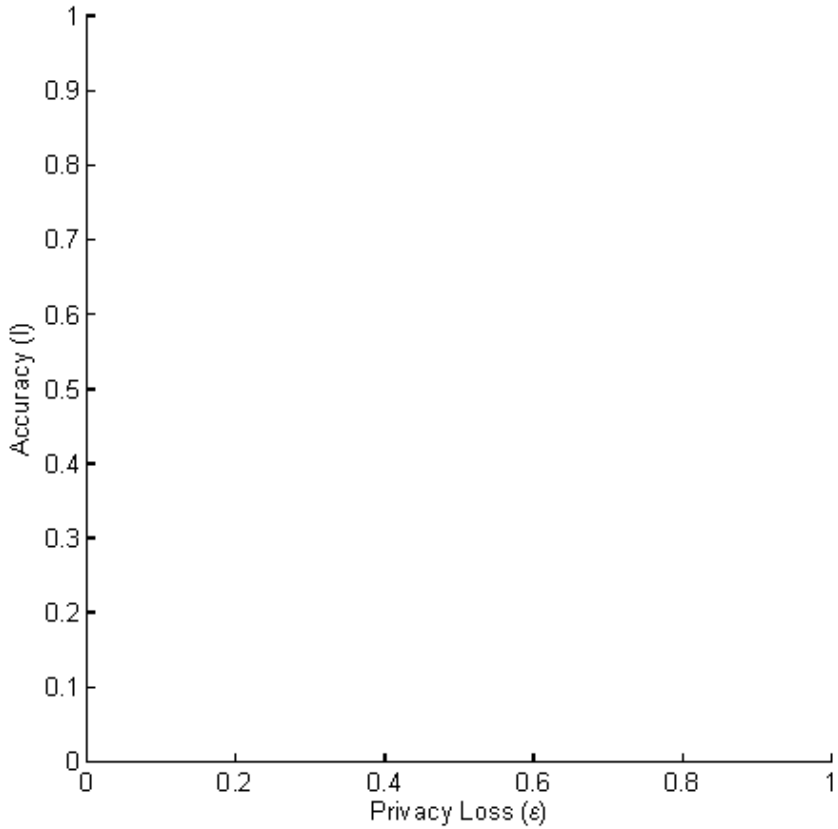

Social Welfare Maximization

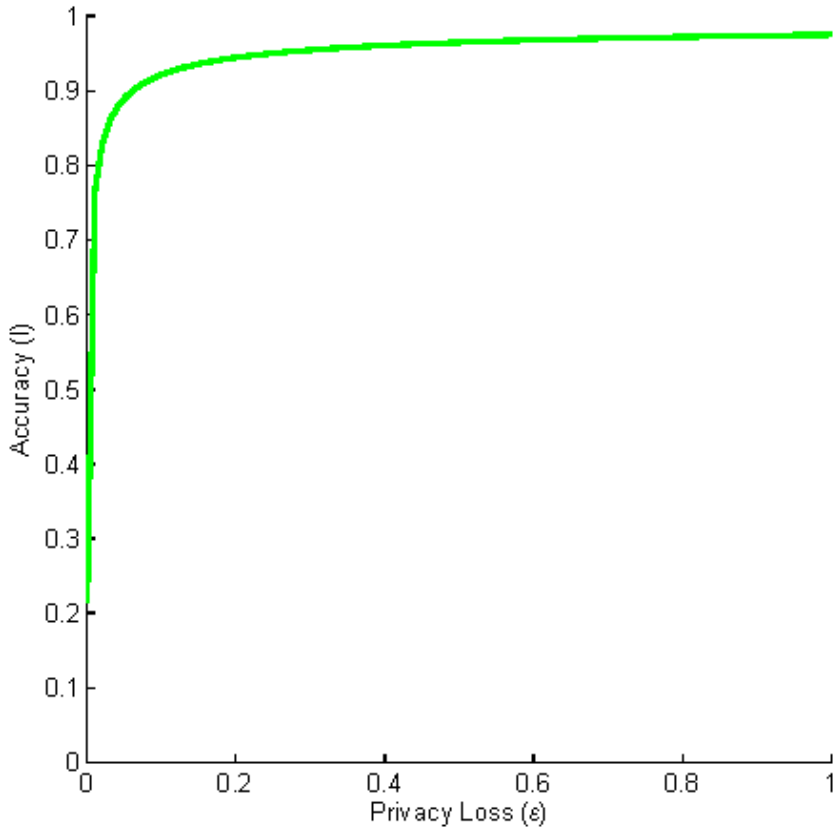

Social Welfare Maximization

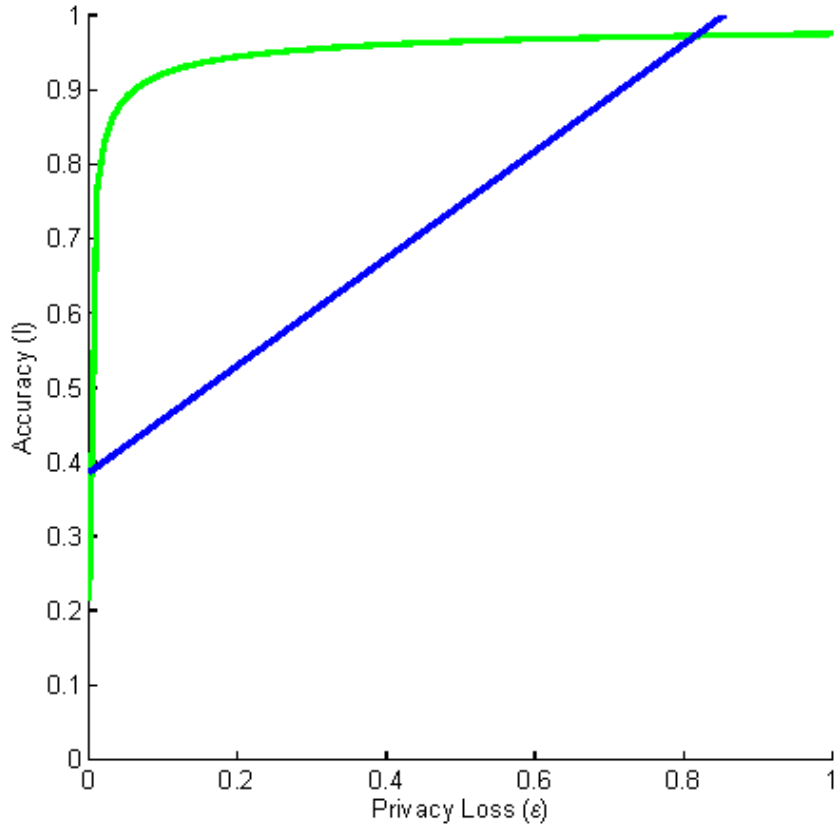

Social Welfare Maximization

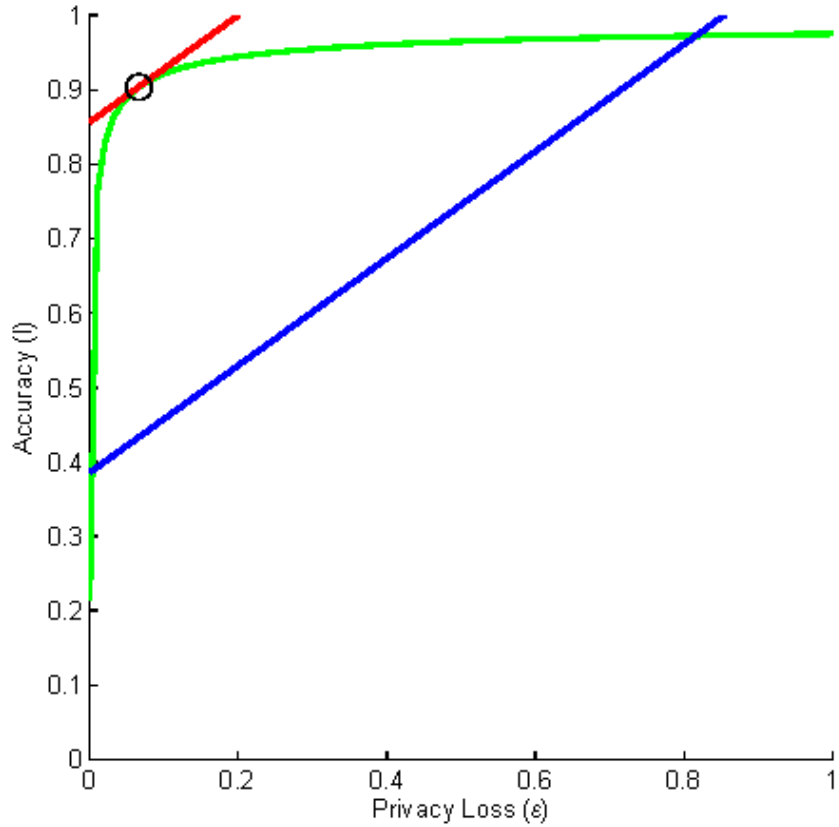

Social Welfare Maximization

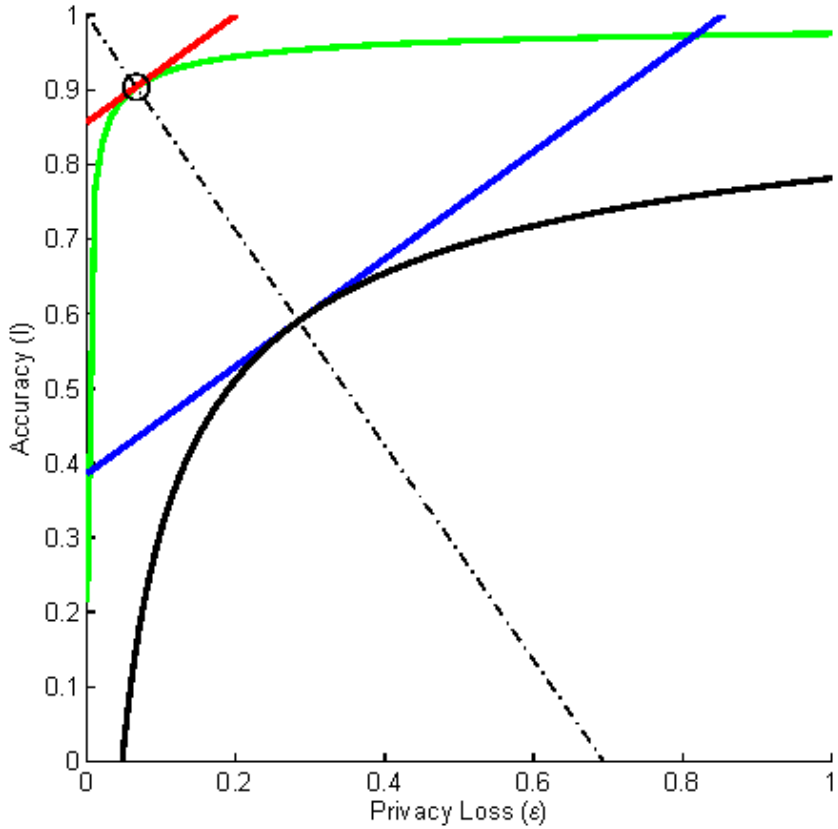

A Little Computer Science

- Formally private data publication systems (e.g., differential privacy, Dwork et al. 2006, but 2017 is much clearer)
- Database definition, including neighboring databases
- Query sets
- Randomized query response mechanism
- Formal privacy definition
- Measure of release data accuracy

Databases and Neighbors

- Basic databases are single tables with rows representing entities and columns representing fields/variables
- Confine attention today to databases where the database schema is limited to finite, discrete outcomes for each field
- Use the histogram representation
  - Each row of the histogram is an image of a row in the database schema augmented with its count in the database
  - Statisticians call this the interior cells of the complete contingency table
- Neighboring databases are those for which all the database constraints are satisfied with the minimum number of changes to the records

Query Sets

- Query sets specify the allowable functions relating the database to an answer space
- Confine attention today to counting queries on the histogram representation of the database

Randomized Query Response Mechanism

- This is the heart of differential privacy
- Randomized query response mechanism add noise to the query responses before releasing the query answer
- The database reconstruction theorem (Dinur and Nissim 2003) establishes that randomized query response mechanisms are necessary, but not sufficient, to prevent arbitrarily accurate reconstruction of the confidential database from a sequence of counting queries

Formal Privacy Definition

- The formal privacy definition puts conditions on the randomized query response mechanism that globally bound the information leakage according to a parameter ε, that is usually called the privacy-loss budget
- We will use the basic differential privacy definition today
- Differential privacy bounds the Bayes factor associated with the worst-case inferential disclosure for all neighboring databases of all databases consistent with the schema

Accuracy of Released Data

- Since randomized query response mechanisms add noise to the correct answer from the database, an accuracy measure is required to compare the protected answer to the true answer
- There are many accuracy measures that might be suitable
- Most depend on the absolute difference between the true answer and the released answer
- We will confine attention today to two accuracy measures that have this property:
  - Normalized total variation distance (basically L_1 distance)
  - Statistical precision relative to the precision in the confidential data (basically L_2 distance)

Example 1: Randomized Response

- Randomized response is differentially private
- Privacy loss is bounded by the maximum Bayes factor

$$\max BF = \frac{\frac{Pr[SQ = Yes|A = Yes]}{Pr[SQ = No|A = Yes]}}{\frac{Pr[SQ = Yes]}{Pr[SQ = No]}} = \frac{Pr[A = Yes|SQ = Yes]}{Pr[A = Yes|SQ = No]} = \frac{(1/2) + (1 - 1/2)\,1/2}{(1 - 1/2)\,1/2} = 3$$

- Bound is the logarithm of the maximum Bayes factor
- If
  - Sensitive question asked with probability ½
  - And innocuous question is "yes" with probability ½
  - Then the maximum Bayes factor is 3, and ln 3 = 1.1
- The privacy-loss expenditure (ε-differential privacy) is 1.1
- Sources: Warner (1965) [link] and Greenberg, Abdel-Latif, Simmons, and Horvitz (1969) [link]. SDL uses: Fienberg and Steele (1998) [link], Du and Zhan (2003) [link] and Erlingsson, Vasyl and Korolova (2014) [link].

What Happens to Data Quality?

- Use relative sampling precision

$$Rel.Precision = \frac{\{Pr[Ask\ Sensitive\ Q]\}^2 \frac{n}{\theta(1-\theta)}}{\frac{n}{\theta(1-\theta)}} = \left\{\frac{1}{2}\right\}^2 = 0.25$$

- If
  - Privacy loss is ln 3
  - Then, relative sampling precision is 25% of the most accurate estimator

Production Possibilities Frontier/Risk-Utility/Receiver Operation Characteristics for Statistical Disclosure Limitation via Randomized Response

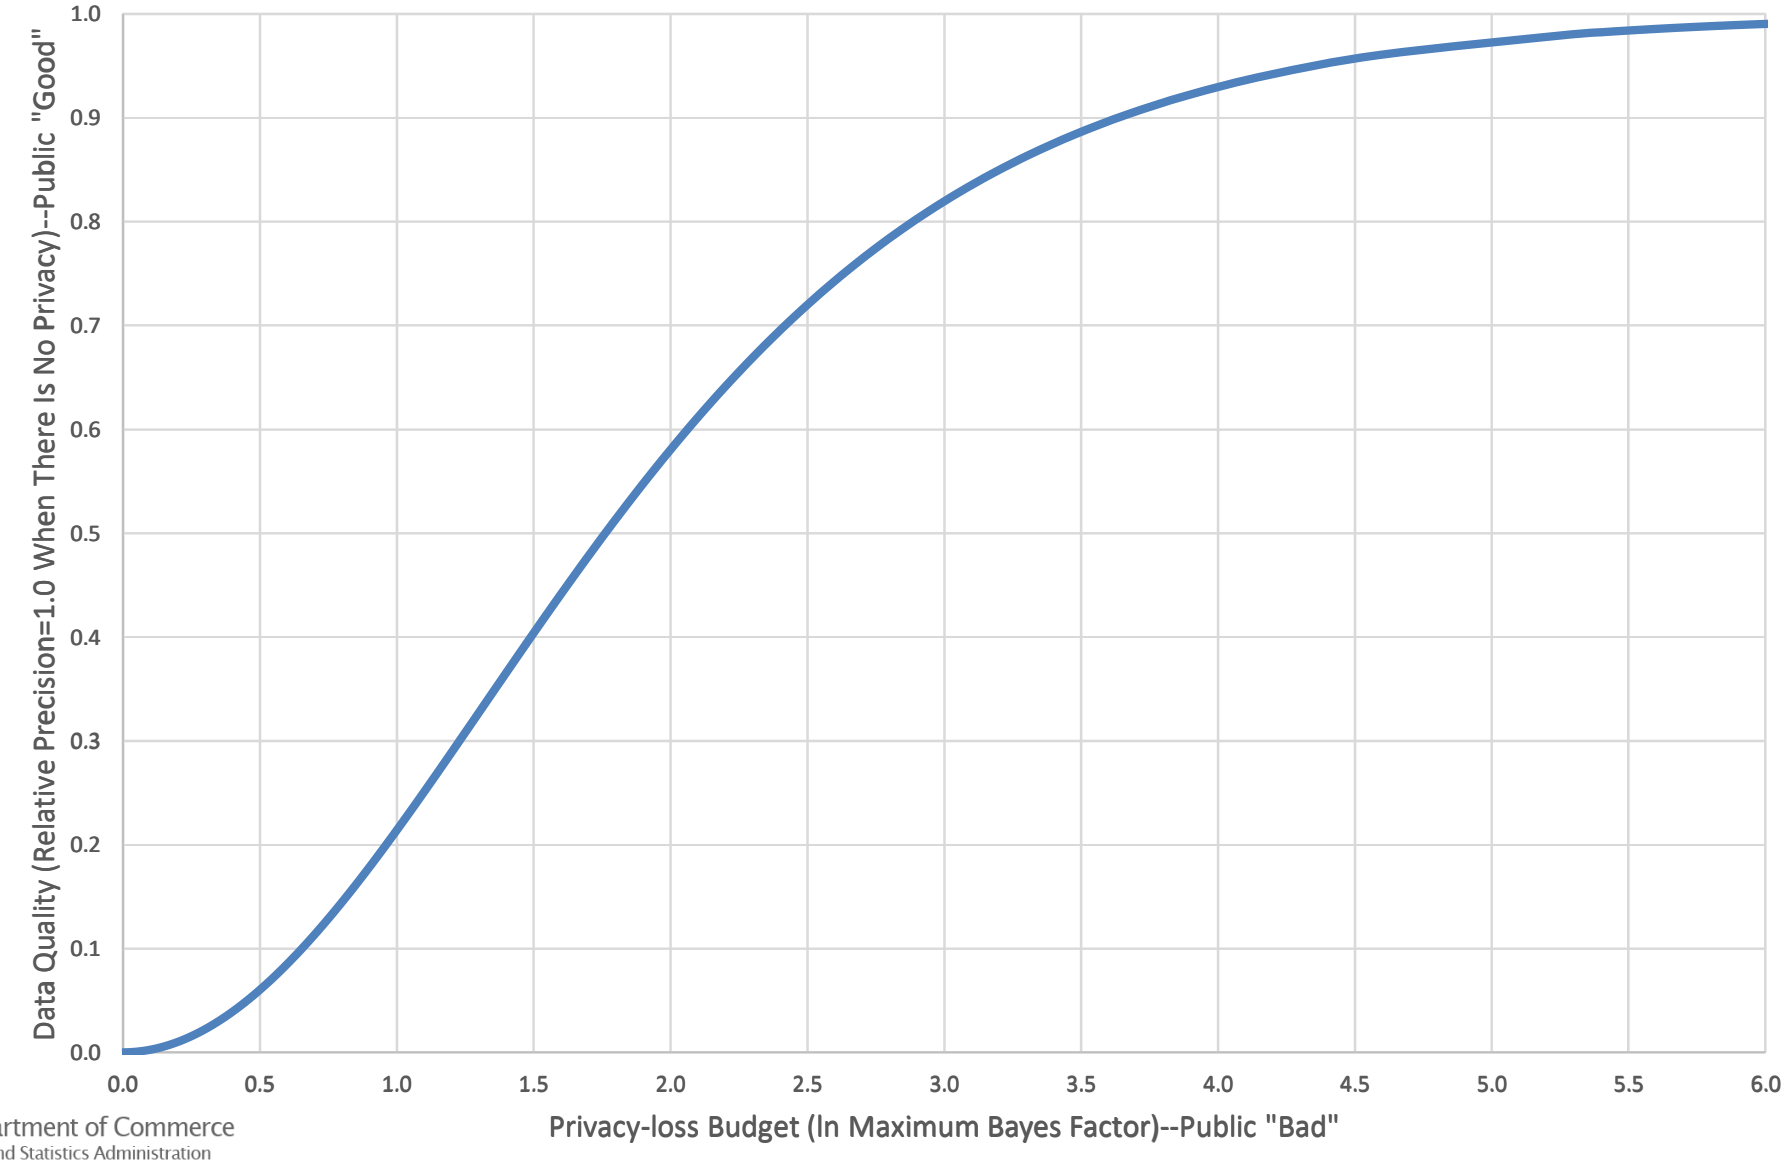

Example 2: Laplace Mechanism for Simple Tables

- The simple tables in this example are based on the 2010 Census of Population PL94-171 release, known popularly as the redistricting data
- These data are used to redraw every legislative district from Congressional all the way down to village councils in every state every ten years
- They must be released by April 1st of the year following a decennial census (for the 2010 Census, April 1, 2011)

Database

- The database is the single-table image of the microdata with one record per entity
- In this presentation think: person records for occupied housing units

Sample Data Table

PID	Tract	Block	GeoID	White	Black	Hispanic	Voting Age	NH-White Alone	NH-Black Alone	NH-Both	H-White Alone	H-Black Alone	H-Both
R01	1	1	11	1	0	1	1	0	0	0	1	0	0
R02	1	1	11	1	0	0	1	1	0	0	0	0	0
R03	1	1	11	1	0	0	1	1	0	0	0	0	0
R04	1	1	11	1	0	0	1	1	0	0	0	0	0
R05	1	1	11	0	1	0	1	0	1	0	0	0	0
R06	1	2	12	0	1	0	1	0	1	0	0	0	0
R07	1	2	12	0	1	1	1	0	0	0	0	1	0
R08	2	1	21	1	0	0	1	1	0	0	0	0	0
R09	2	1	21	1	0	0	0	1	0	0	0	0	0
R10	2	1	21	1	0	0	1	1	0	0	0	0	0
R11	2	2	22	1	0	0	0	1	0	0	0	0	0
R12	2	2	22	1	0	0	1	1	0	0	0	0	0
R13	2	2	22	1	0	0	1	1	0	0	0	0	0
R14	2	2	22	1	0	0	1	1	0	0	0	0	0
R15	2	2	22	1	0	0	1	1	0	0	0	0	0
R16	2	2	22	1	0	0	1	1	0	0	0	0	0
R17	2	2	22	1	1	0	1	0	0	1	0	0	0
R18	2	3	23	1	1	1	0	0	0	0	0	0	1
R19	2	3	23	1	0	0	0	1	0	0	0	0	0
R20	3	1	31	1	0	0	0	1	0	0	0	0	0
TOTAL			20	17	5	3	15	14	2	1	1	1	1

This table is the person records only from occupied households (simulated data).

Database Schema

- To statisticians: sample space
- The legal combinations of each of the tabulation variables in the database
- Image of every legal record in the database
- Structural zeros are imposed by deleting rows in this image
- Other constraints (e.g., total population counts, group quarters types) are represented by linear equalities and inequalities

Sample Database Schema

SchemaID	GeoID	White	Black	Hispanic	Voting Age	NH-White Alone	NH-Black Alone	NH-Both	H-White Alone	H-Black Alone	H-Both	Sensitivity 1	Sensitivity 2
1	11	0	1	0	0	0	1	0	0	0	0	1	1
2	11	0	1	0	1	0	1	0	0	0	0	1	1
3	11	0	1	1	0	0	0	0	0	1	0	1	1
4	11	0	1	1	1	0	0	0	0	1	0	1	1
5	11	1	0	0	0	1	0	0	0	0	0	1	1
6	11	1	0	0	1	1	0	0	0	0	0	1	1
7	11	1	0	1	0	0	0	0	1	0	0	1	1
8	11	1	0	1	1	0	0	0	1	0	0	1	1
9	11	1	1	0	0	0	0	1	0	0	0	1	2
10	11	1	1	0	1	0	0	1	0	0	0	1	2
11	11	1	1	1	0	0	0	0	0	0	1	1	2
12	11	1	1	1	1	0	0	0	0	0	1	1	2
13	12	0	1	0	0	0	1	0	0	0	0	1	1
14	12	0	1	0	1	0	1	0	0	0	0	1	1
15	12	0	1	1	0	0	0	0	0	1	0	1	1
16	12	0	1	1	1	0	0	0	0	1	0	1	1
17	12	1	0	0	0	1	0	0	0	0	0	1	1
18	12	1	0	0	1	1	0	0	0	0	0	1	1
19	12	1	0	1	0	0	0	0	1	0	0	1	1
20	12	1	0	1	1	0	0	0	1	0	0	1	1
21	12	1	1	0	0	0	0	1	0	0	0	1	2
22	12	1	1	0	1	0	0	1	0	0	0	1	2
23	12	1	1	1	0	0	0	0	0	0	1	1	2
24	12	1	1	1	1	0	0	0	0	0	1	1	2
25	21	0											
26	21	0											

This is a portion of the database schema that applies to the example database. The full schema has 72 rows. Note that the structural zeros implied by requiring at least one race (either white or black in this example) are imposed. Current OMB standards also imposed.

Neighboring Database

- Neighboring databases are those that differ from the actual database by changing no more than two rows
- This definition occurs because the total population in any geography is not allowed to change (a constraint)
- In this case the linear constraint that total population is exact in every block means that neighboring databases must have the same number of rows in each block

Queries and Sensitivity

- To statisticians: any well-defined summary statistic computed on the database is a query
- We restrict attention to counting queries, which everyone understands the same way
- Sensitivity is the maximum amount any allowable query can change (in absolute value) when all neighbors of all possible databases are considered

Queries and Sensitivity

- The sensitivity of the query workload for the redistricting data is two
- If the block population can't change, then two neighboring databases must have a difference in value for exactly two persons
- Population count: 10
- White: 5, Black: 5 and White: 4, Black: 6 are neighboring databases
- Abs(5-4)+Abs(5-6) = 2

Randomized Query Mechanism

- To statisticians: input or output noise infusion as the disclosure limitation technology
- Input noise infusion example: randomized response mechanism
- Output noise infusion example: Laplace mechanism
  - Add independent Laplace noise to each query
  - Post process to satisfy nonnegativity and integer constraints
  - We actually use a discrete version of this called the geometric mechanism

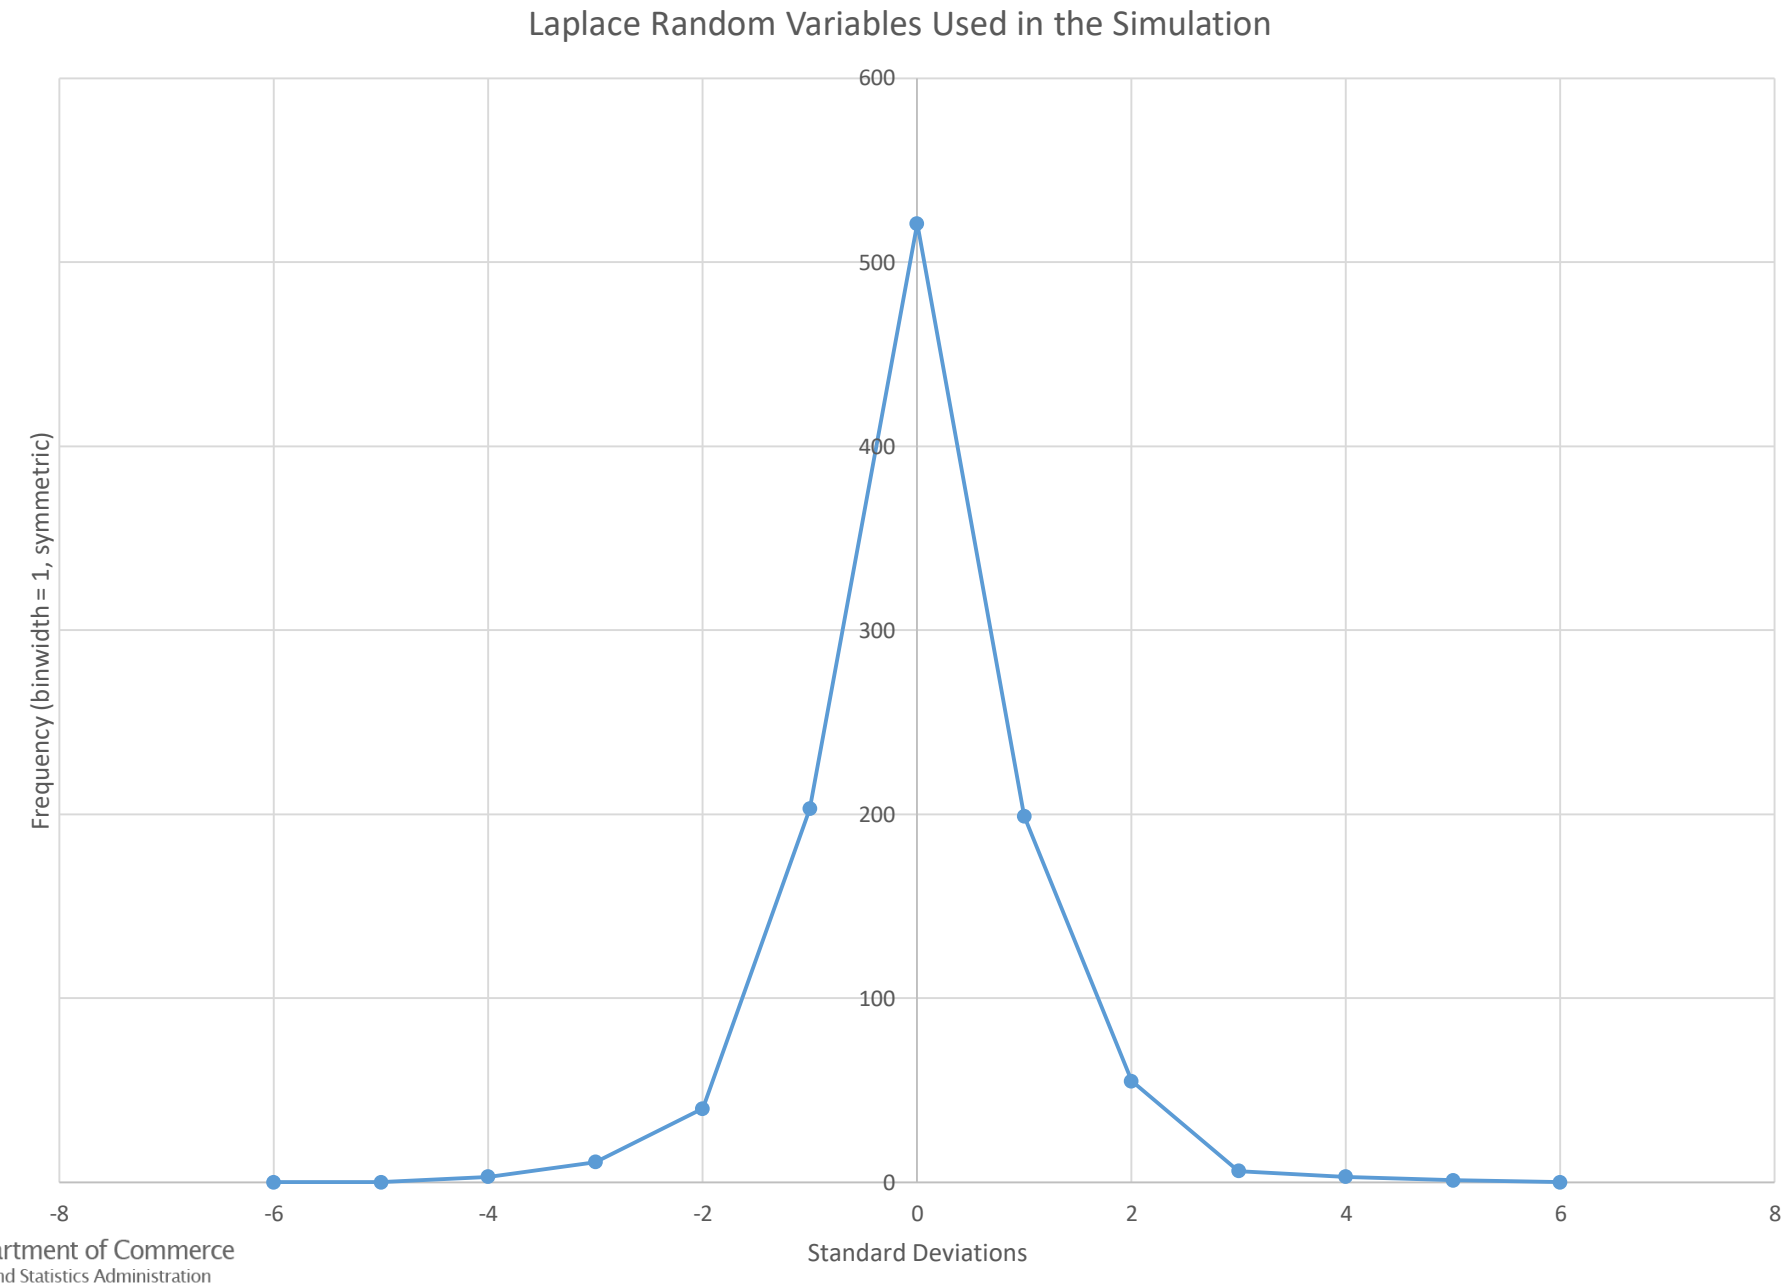
Laplace Random Variables Used in the Simulation
Frequency (binwidth = 1, symmetric)
600
500
400
300
200
100
0
-8
-6
-4
-2
0
2
4
6
8
Standard Deviations

Differential Privacy

- To statisticians: a global bound on the maximum Bayes factor associated with any exact identity or attribute disclosure
- The bound, called ε in differential privacy, is taken over all possible databases with respect to all possible neighbors of those databases
- The “neighbor” definition makes exact identity and attribute disclosures precise

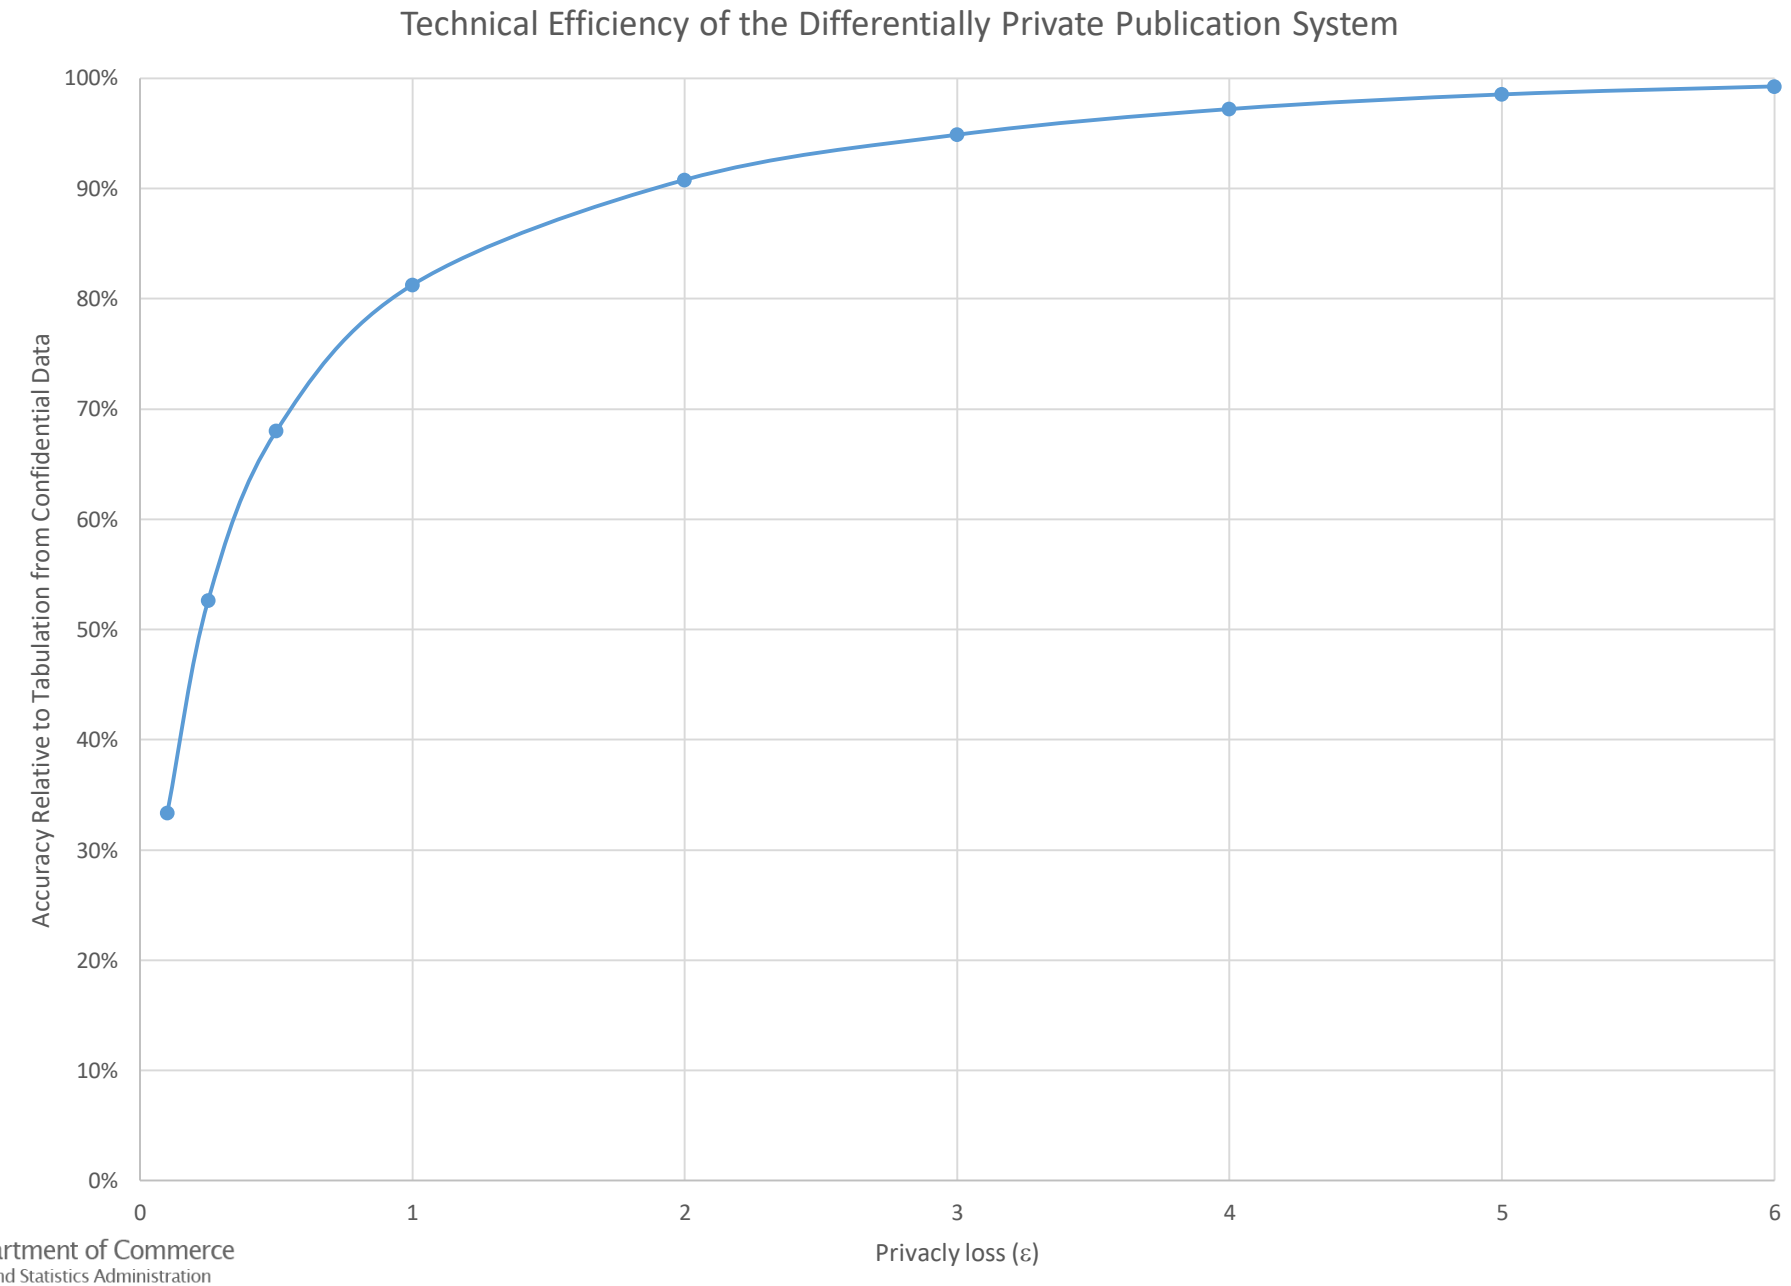

This is the plot of the actual tradeoff between data accuracy (relative to tabulating the confidential data) for the Laplace mechanism as applied to the national table from the example.

Disclosure Limitation is Technology

- The price of increasing data quality (public “good”) in terms of increased privacy loss (public “bad”) is the slope of the technology frontier:
  - Economics: Production Possibilities Frontier (Risk-Return in finance)
  - Forecasting models: Receiver Operating Characteristics Curve
  - Statistical Disclosure Limitation: Risk-Utility Curve (with risk on the x-axis)
- All exactly the same thing
- None able to select an optimal point

Production Possibilities Frontier/Risk-Utility/Receiver Operation Characteristics for Statistical Disclosure Limitation via Randomized Response

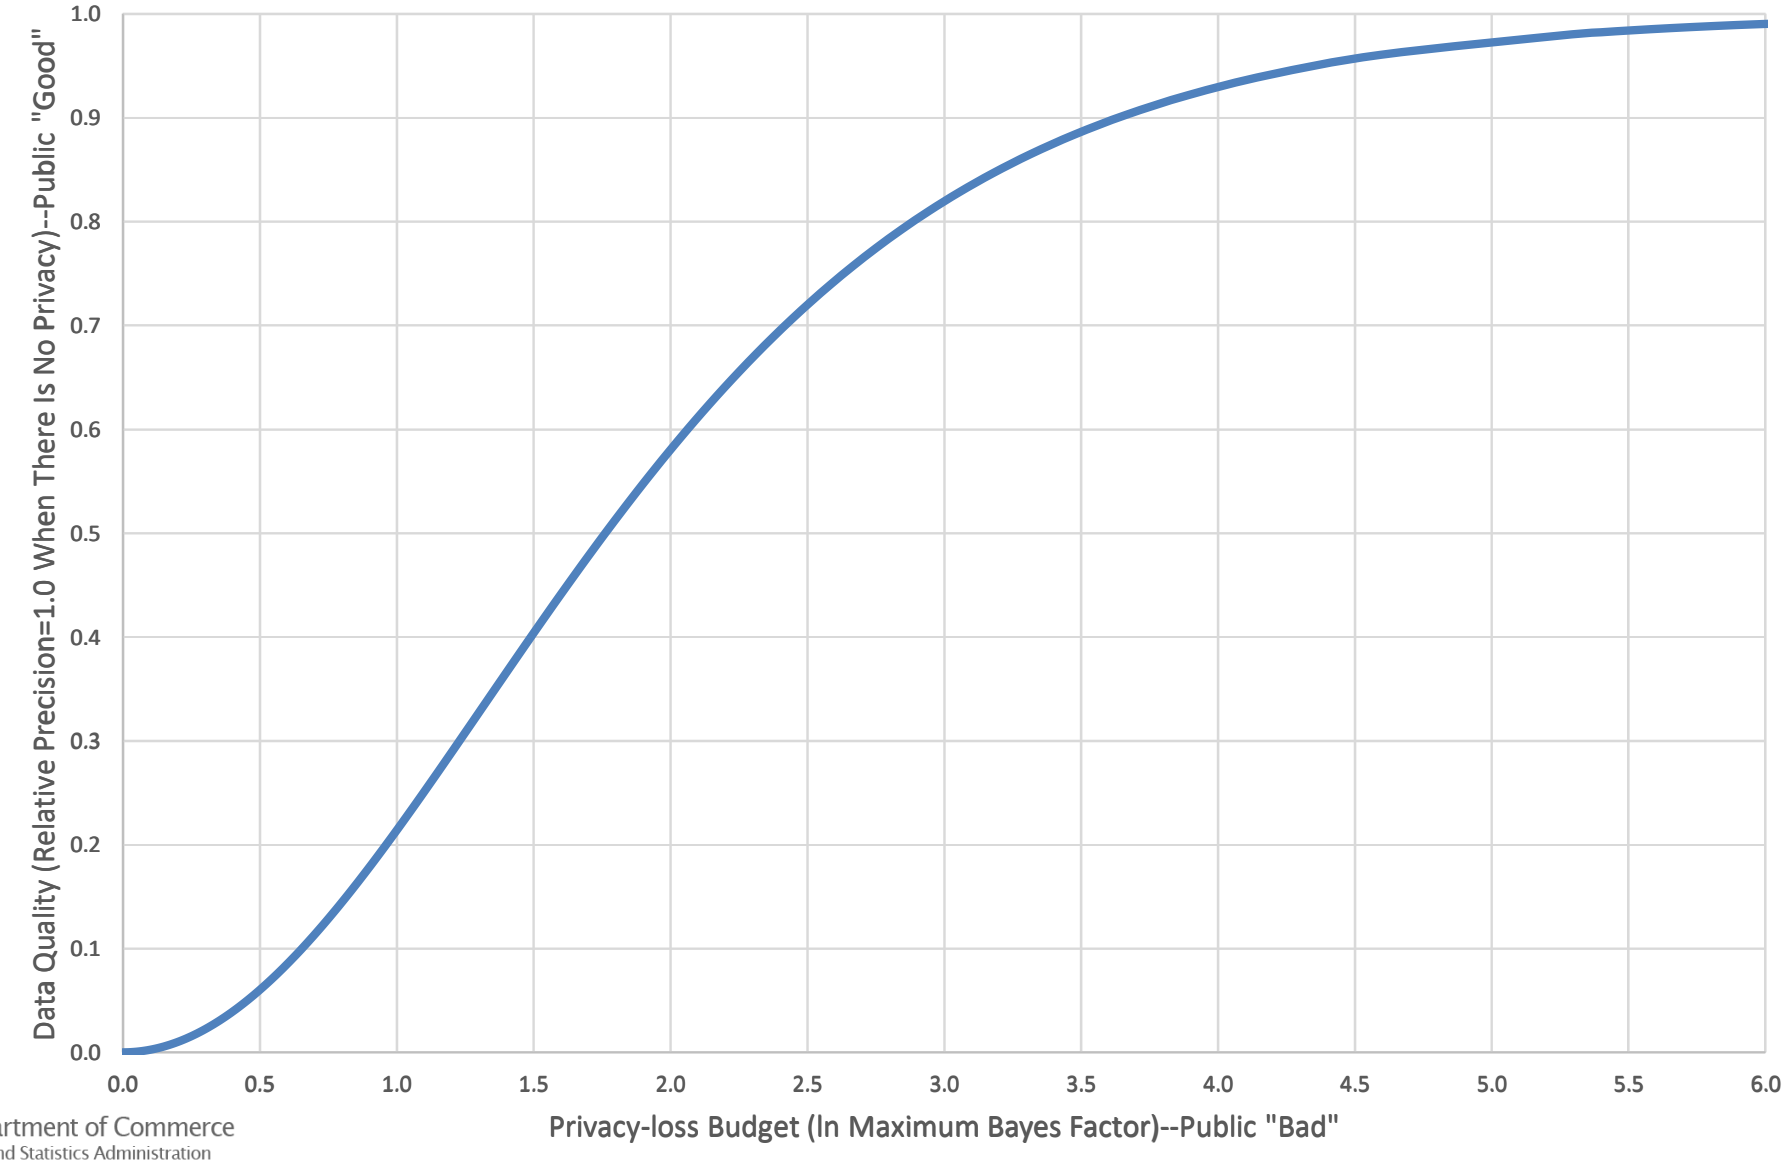

How Do You Set ε?

- Dwork (2008): “The parameter ε in Definition 1 is public. The choice of ε is essentially a social question and is beyond the scope of this paper.” [link, p. 3]
- Dwork (2011): “The parameter ε is public, and its selection is a social question. We tend to think of ε as, say, 0.01, 0.1, or in some cases, ln 2 or ln 3.” [link, p. 91]
- In OnTheMap, ε = 8.9, was required to produce tract-level estimates with acceptable accuracy
- All these settings are differentially private, but they very different global disclosure risk

Production Possibilities Frontier/Risk-Utility/Receiver Operation Characteristics for Statistical Disclosure Limitation via Randomized Response

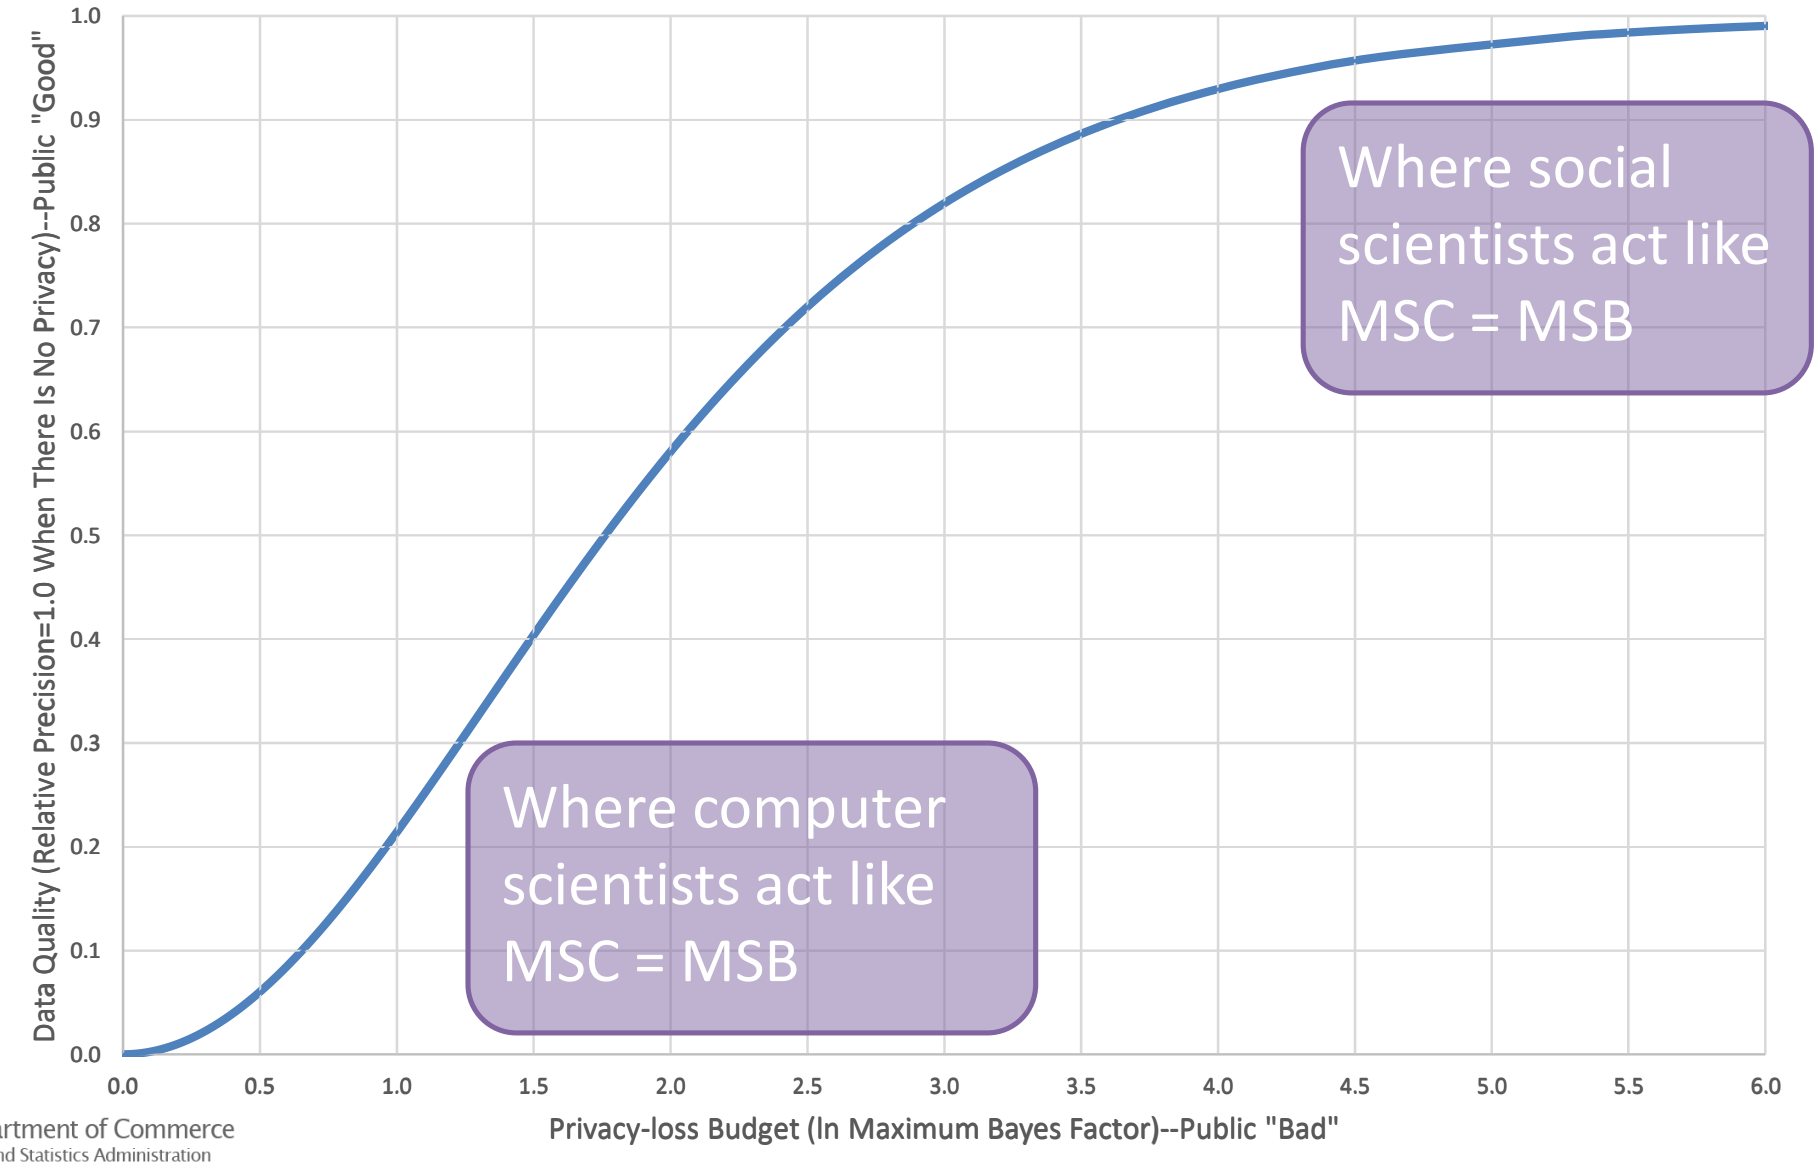

How to Think about the Social Choice Problem

- The marginal social benefit is the sum of all persons' willingness-to-pay for data accuracy with increased privacy loss
- Ian and I estimated the choice parameters from survey data
- The next slide shows an example for randomized response
- This is the problem being addressed by Google in RAPPOR and Apple in iOS 11

See Abowd and Schmutte (2017) [link].

Production Possibilities Frontier/Risk-Utility/Receiver Operation Characteristics for Statistical Disclosure Limitation via Randomized Response

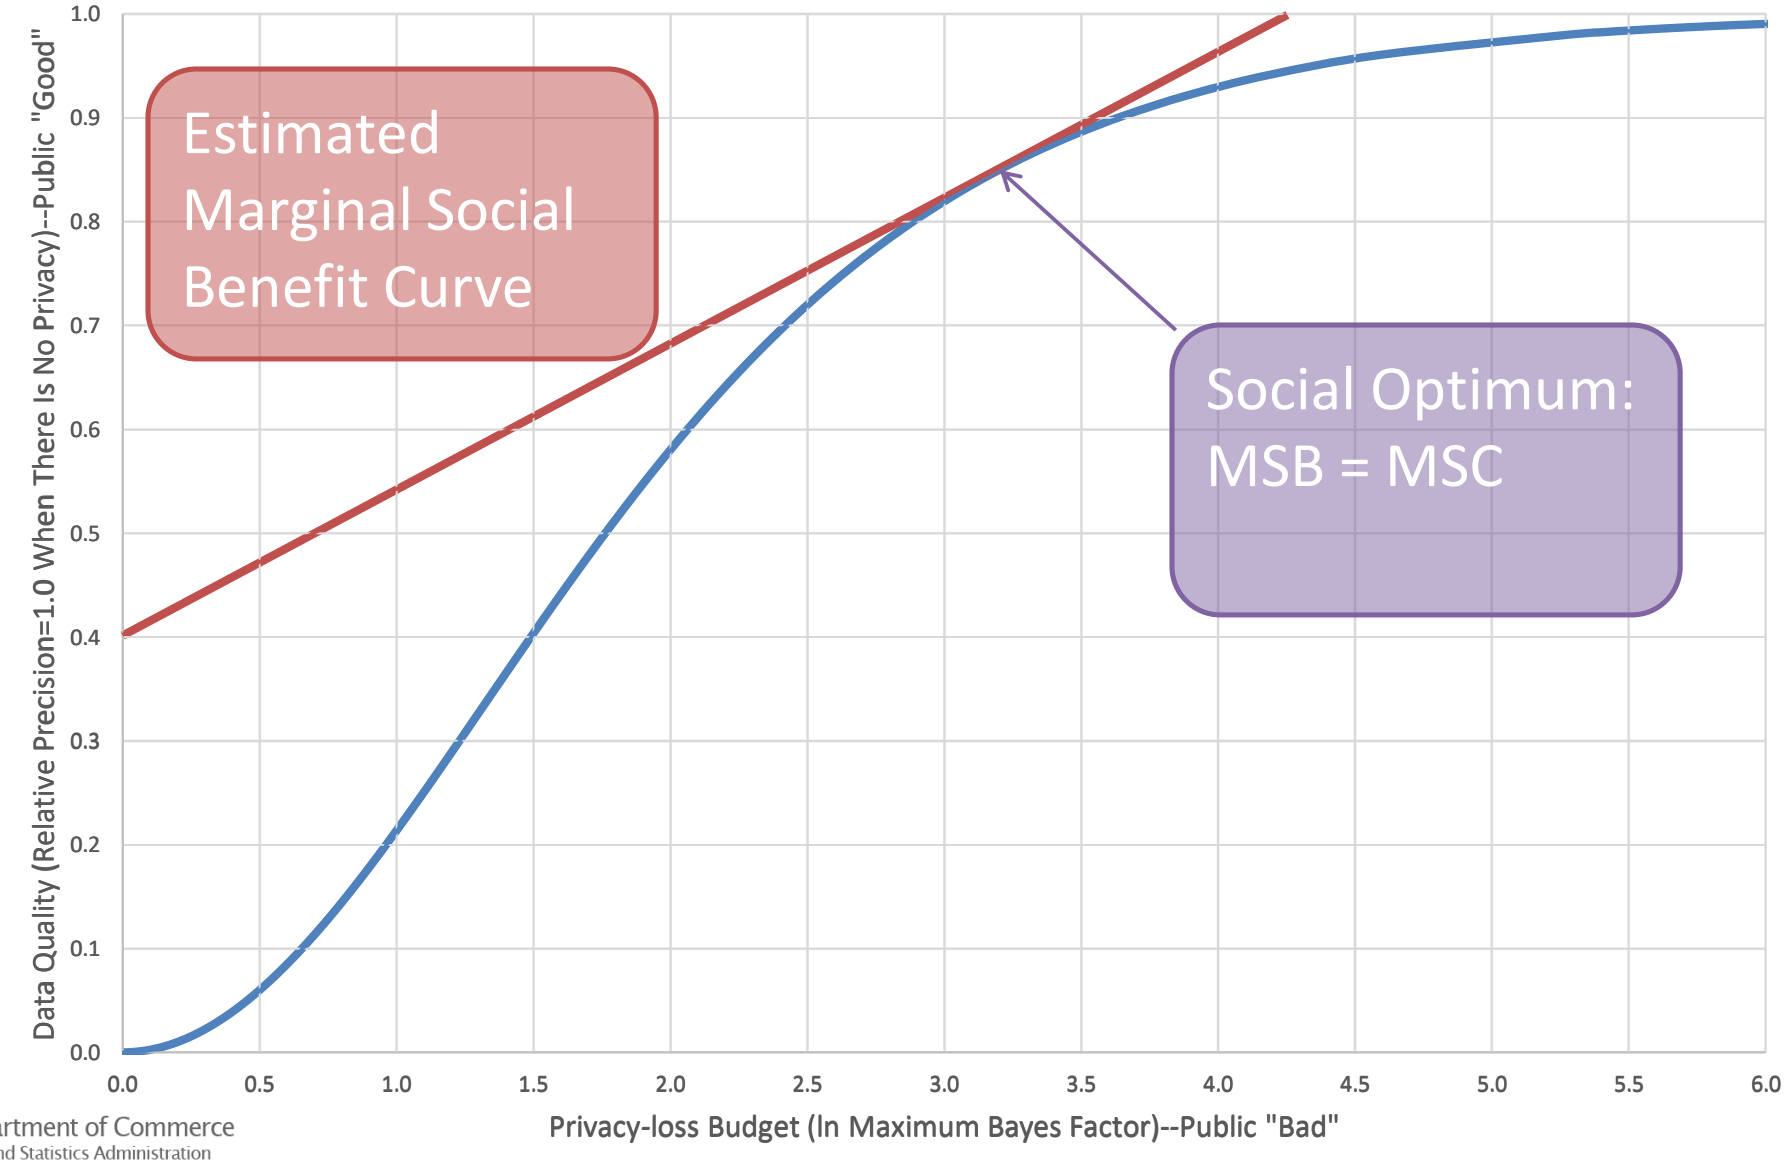

But the Choice Problem for Example 2 is More Challenging

- In the redistricting application, the fitness-for-use is based on
  - Supreme Court one-person one-vote decision
  - Voting Rights Act: requires majority minority districts at all levels, when certain criteria
- The privacy interest is based on
  - Title 13 requirement not to publish exact identifying information

Production Possibilities Frontier/Risk-Utility/Receiver Operation Characteristics for Statistical Disclosure Limitation via Alternative Mechanisms

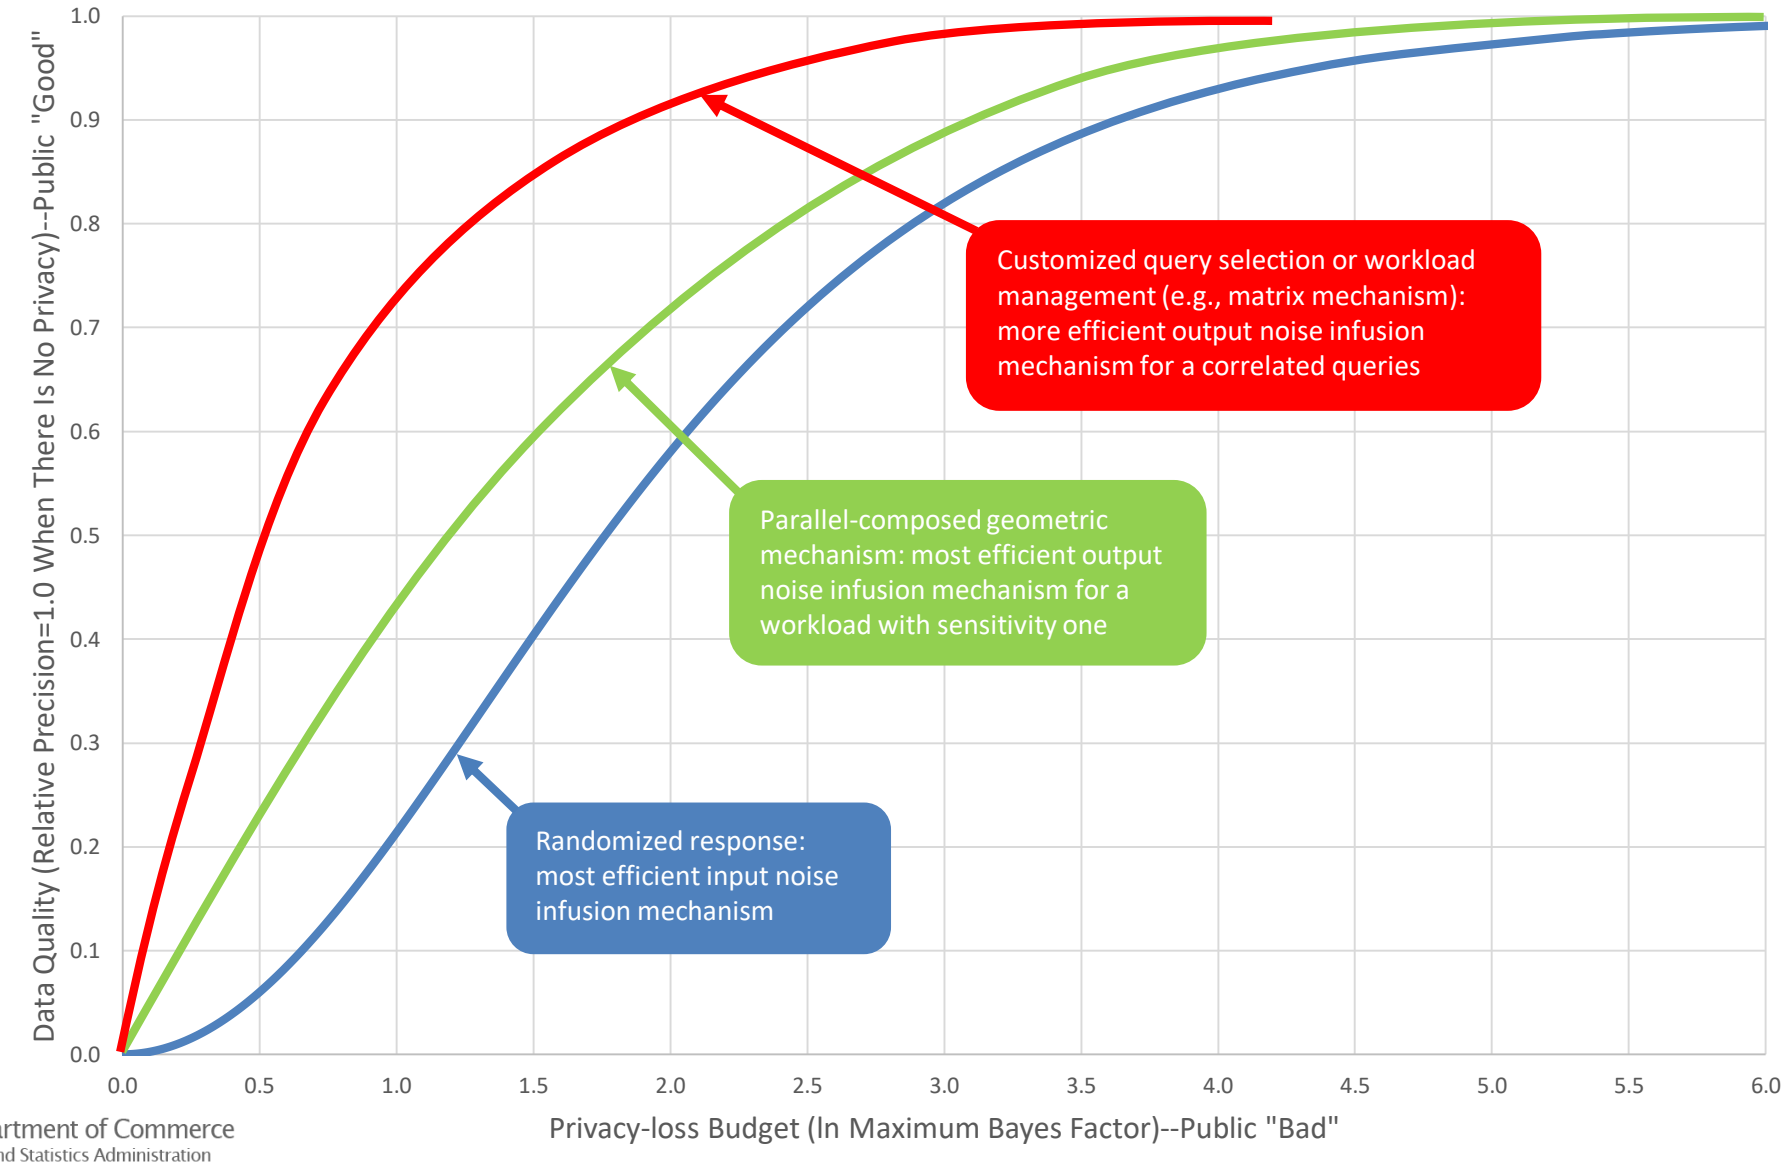

Production Possibilities Frontier/Risk-Utility/Receiver Operation Characteristics for Statistical Disclosure Limitation via Alternative Mechanisms

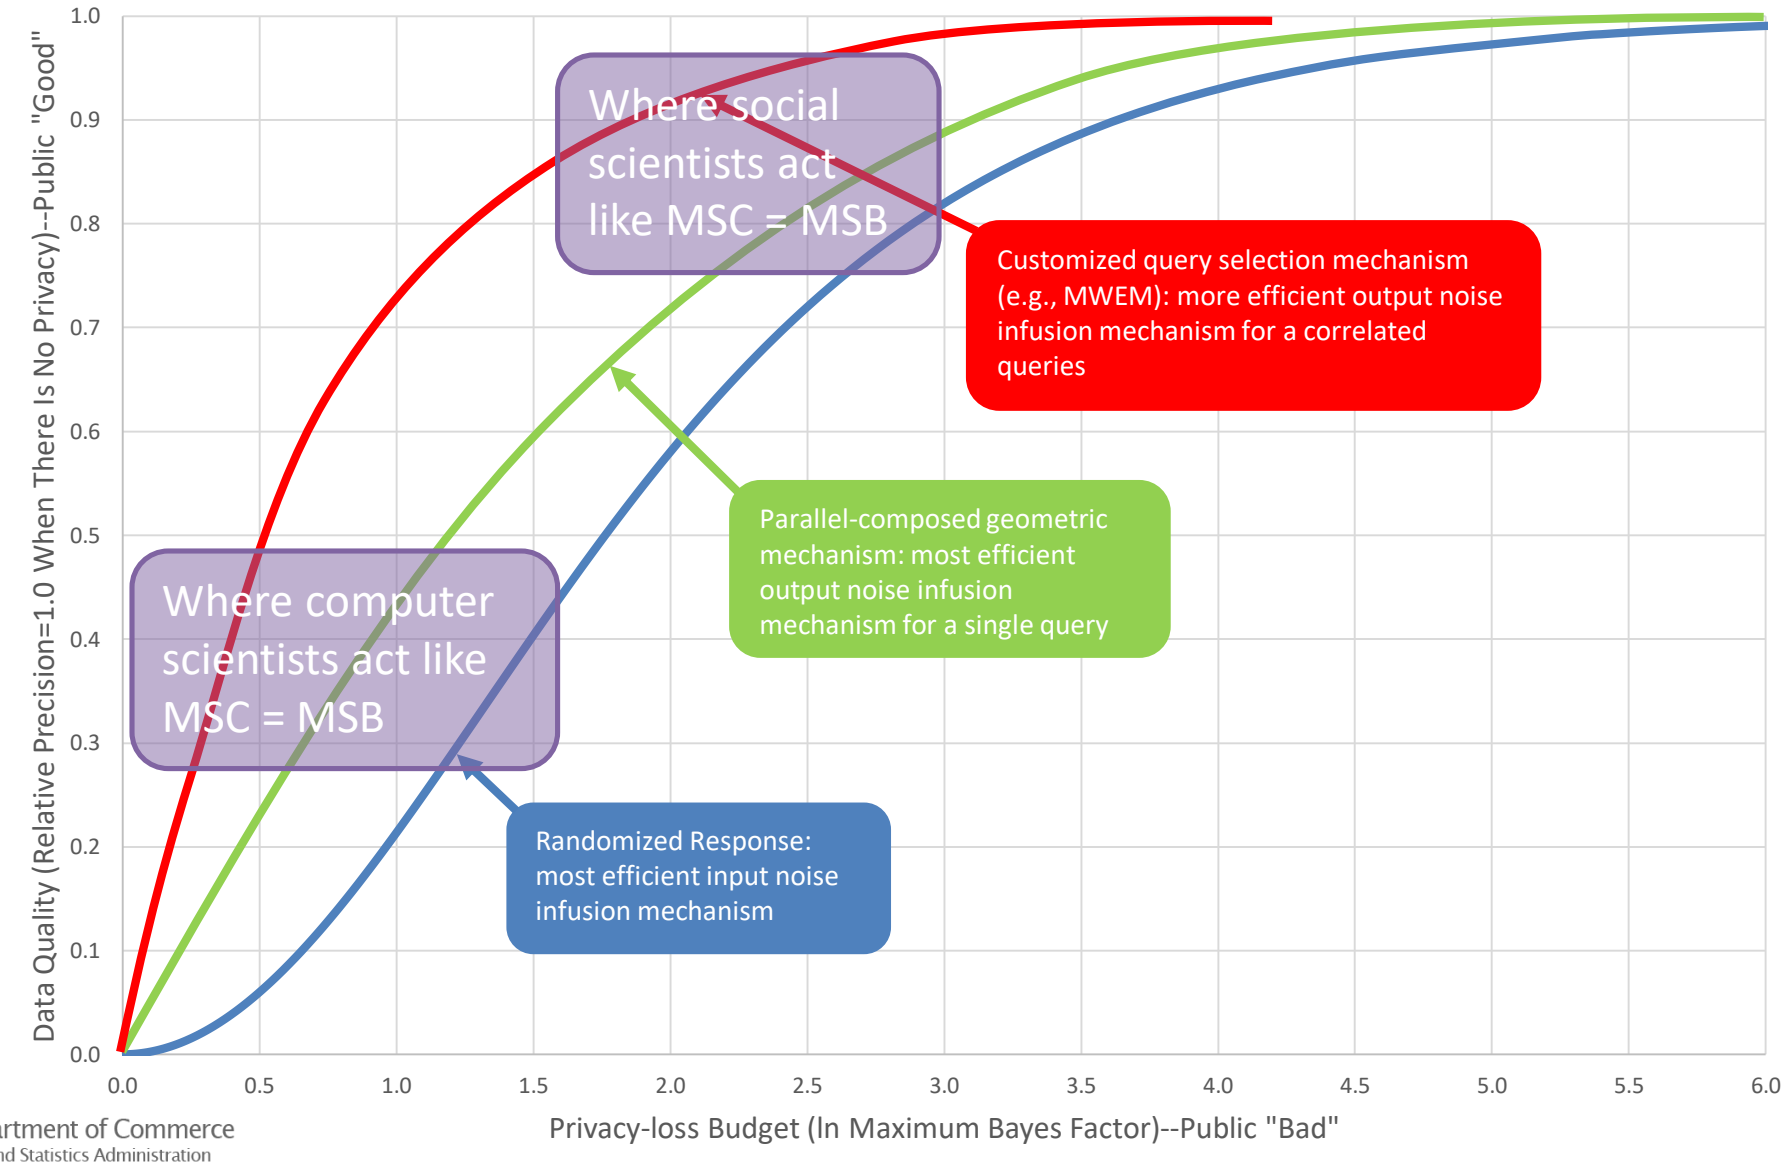

Production Possibilities Frontier/Risk-Utility/Receiver Operation Characteristics for Statistical Disclosure Limitation via Alternative Mechanisms

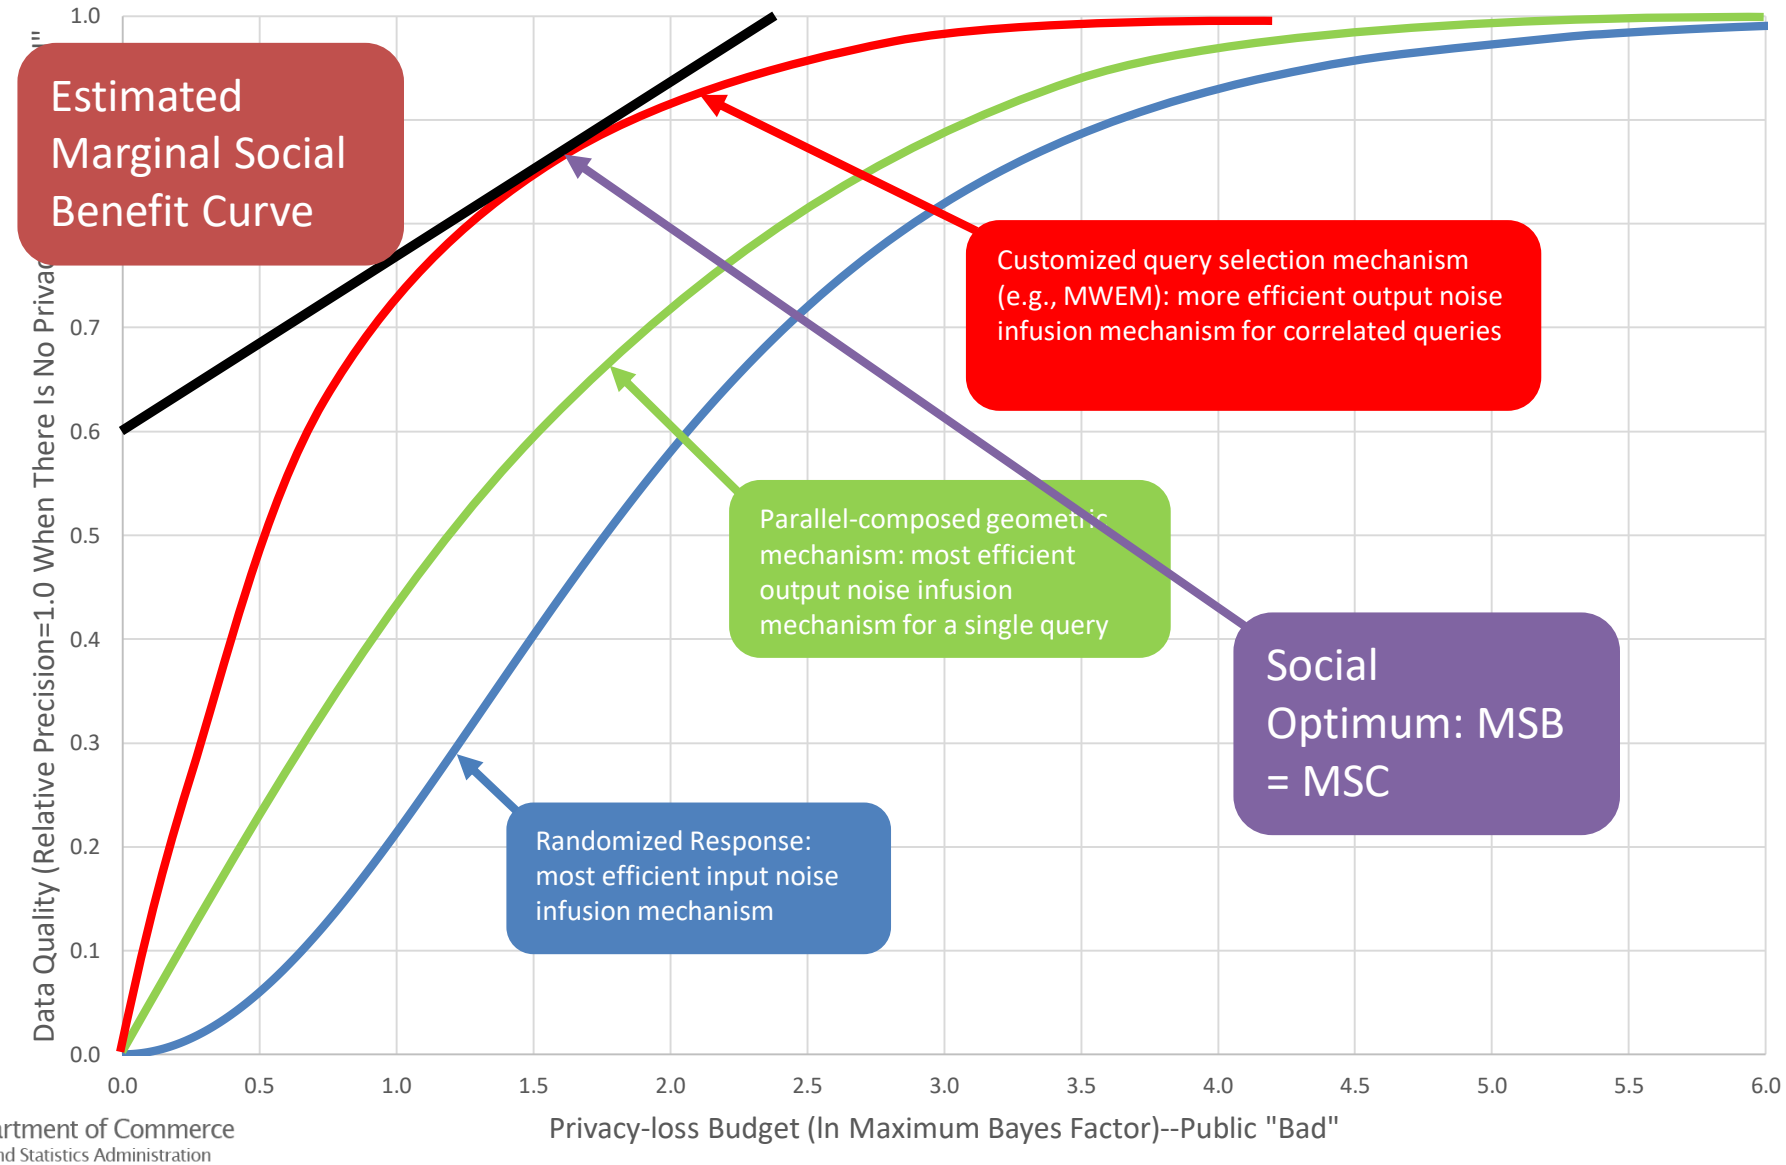

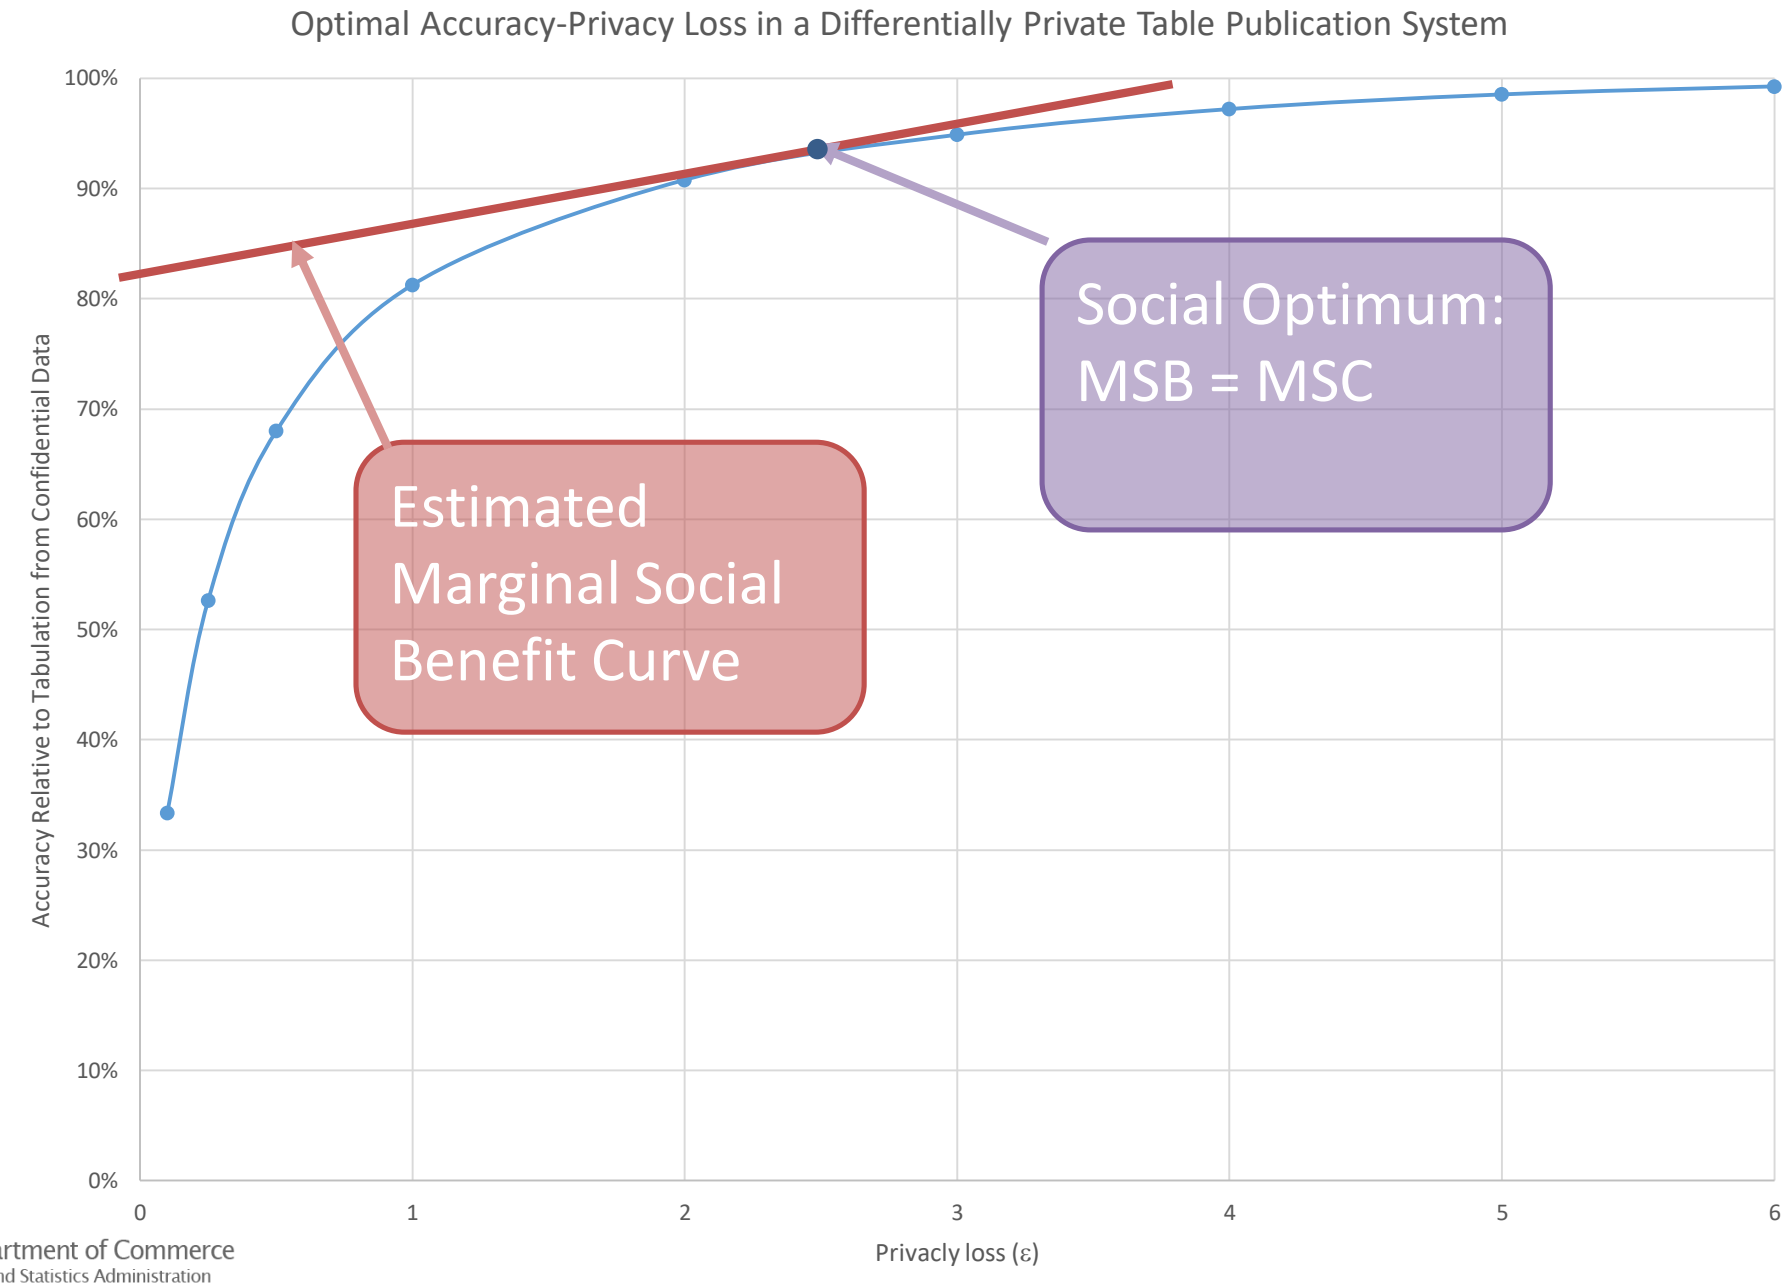
Optimal Accuracy-Privacy Loss in a Differentially Private Table Publication System
Accuracy Relative to Tabulation from Confidential Data
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
0
1
2
3
4
5
6
Privacly loss (ε)
Social Optimum:
MSB = MSC
Estimated
Marginal Social
Benefit Curve

Takeaways

- This is new ground for official statisticians: they did not have to think explicitly about the social choice problem of competing interests on accuracy and privacy using traditional disclosure avoidance methods
- But, it is not foreign territory, since the invention of sampling theory, official statisticians have thought about the tradeoff between accuracy and all aspects of the design of surveys

Thank you.

john.maron.abowd@census.gov

Selected References

- Dinur, Irit and Kobbi Nissim. 2003. Revealing information while preserving privacy. In *Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*(PODS '03). ACM, New York, NY, USA, 202-210. DOI: 10.1145/773153.773173.
- Dwork, Cynthia, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. in Halevi, S. & Rabin, T. *(Eds.)* Calibrating Noise to Sensitivity in Private Data Analysis *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings, Springer Berlin Heidelberg*, 265-284, DOI: 10.1007/11681878_14.
- Dwork, Cynthia. 2006. Differential Privacy, *33rd International Colloquium on Automata, Languages and Programming, part II (ICALP 2006), Springer Verlag, 4052*, 1-12, ISBN: 3-540-35907-9.
- Dwork, Cynthia and Aaron Roth. 2014. *The Algorithmic Foundations of Differential Privacy*. Foundations and Trends in Theoretical Computer Science. Vol. 9, Nos. 3–4. 211–407, DOI: 10.1561/0400000042. [download]
- Dwork, Cynthia. Undated. The State of the Art. Slide presentation. [download]
- Dwork, Cynthia, Frank McSherry and Kunal Talwar. 2007. The price of privacy and the limits of LP decoding. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*(STOC '07). ACM, New York, NY, USA, 85-94. DOI:10.1145/1250790.1250804.
- Muthukrishnan, S. and Aleksandar Nikolov. 2012. Optimal Private Halfspace Counting via Discrepancy, *CoRR (Computing Research Repository), abs/1203.5453.* [download]
- Kasiviswanathan, Shiva Prasad, Mark Rudelson and Adam Smith. 2012. The Power of Linear Reconstruction Attacks, *CoRR (Computing Research Repository),* abs/1210.2381.
- Dwork, Cynthia, Adam Smith, Thomas Steinke, Jonathan Ullman, and Salil Vadhan. 2015. "Robust Traceability from Trace Amounts." In IEEE Symposium on Foundations of Computer Science (FOCS 2015). Berkeley, California, 10/18-20/2015.
- Machanavajjhala, Ashwin, Daniel Kifer, John M. Abowd , Johannes Gehrke, and Lars Vilhuber. 2008. Privacy: Theory Meets Practice on the Map, International Conference on Data Engineering (ICDE) 2008: 277-286, doi:10.1109/ICDE.2008.4497436.
- Dwork, Cynthia and Moni Naor. 2008. On the Difficulties of Disclosure Prevention in Statistical Databases or The Case for Differential Privacy, *Journal of Privacy and Confidentiality*: Vol. 2: Iss. 1, Article 8. Available at: http://repository.cmu.edu/jpc/vol2/iss1/8.
- Kifer, Daniel and Ashwin Machanavajjhala. 2011. No free lunch in data privacy. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data* (SIGMOD '11). ACM, New York, NY, USA, 193-204. DOI:10.1145/1989323.1989345.
- Beckman, Petr. 1971. *A History of Pi*. Barnes and Noble, New York, NY. ISBN:0-88029-418-3.
- Warner, Stanley. L. 1965. Randomized Response: A Survey Technique for Eliminating Evasive Answer Bias *Journal of the American Statistical Association, 60*, 63-69.
- Greenberg, Bernard G., Abdel-Latif Abul-Ela, Walt R. Simmons, and Daniel G. Horvitz. 1969. The Unrelated Question Randomized Response Model: Theoretical Framework. *Journal of the American Statistical Association, 64*, 520-539.
- Fienberg, Stephen E. Udi E. Makov and Russel J. Steele. 1998. Disclosure Limitation Using Perturbation and Related Methods for Categorical Data. *Journal of Official Statistics,* Vol. 14, No. 4 (December): 485-502.[link]
- Du, Wenliang and Zhijun Zhan. 2003. Using randomized response techniques for privacy-preserving data mining. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining* (KDD '03). ACM, New York, NY, USA, 505-510. DOI:10.1145/956750.956810.
- Erlingsson, Úlfar, Vasyl Pihur and Aleksandra Korolova. 2014. RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security* (CCS '14). ACM, New York, NY, USA, 1054-1067. DOI:10.1145/2660267.2660348.
- Dwork, Cynthia. 2008. Differential Privacy: A Survey of Results. In Agrawal, M., Du, D.; Duan, Z. and Li, A. *(Eds.). Theory and Applications of Models of Computation: 5th International Conference, TAMC 2008, Xi'an, China, April 25-29, 2008. Proceedings, Springer Berlin Heidelberg*, 1-19. [download]
- Dwork, Cynthia. 2011. A firm foundation for private data analysis. *Communications ACM* 54, 1 (January): 86-95. DOI:10.1145/1866739.1866758.
- Abowd, John M. and Ian M. Schmutte. 2017. Revisiting the economics of privacy: Population statistics and confidentiality protection as public goods. Labor Dynamics Institute, Cornell University, Labor Dynamics Institute, Cornell University. [download]
- Wasserstein, Ron L. and Nicole A. Lazar. 2016. The ASA's Statement on p-Values: Context, Process, and Purpose. *The American Statistician, 70*, 129-133. DOI: 10.1080/00031305.2016.1154108.
- Dwork, Cynthia, Feldman V., Hardt M., Pitassi T., Reingold O., Roth A. 2015. The reusable holdout: Preserving validity in adaptive data analysis. Science, 349, 6248 (August 7):636-8. DOI:10.1126/science.aaa9375.
- Apple, Inc. 2016. Apple previews iOS 10, the biggest iOS release ever. Press Release (June 13). URL=http://www.apple.com/newsroom/2016/06/apple-previews-ios-10-biggest-ios-release-ever.html.