Negotiating Privacy-Utility Trade-Offs under Differential Privacy

Gerome Miklau (Tumult Labs and UMass Amherst)

Michael Hay (Tumult Labs and Colgate University)

Ashwin Machanavajjhala (Tumult Labs and Duke University)

Amritha Pai (Tumult Labs)

William Sexton (Tumult Labs)

PEPR '22

JUNE 23, 2022

https://collegescorecard.ed.gov/

BROOKINGS INSTITUTE OP-ED

“This represents a huge step toward transparency in higher education.”

“Parents, students, college leaders, journalists, policy makers, and researchers are now empowered to more **empirically evaluate** thousands of U.S. post-secondary institutions in terms of their **contributions to student economic success.”**

Student earnings, X years after degree

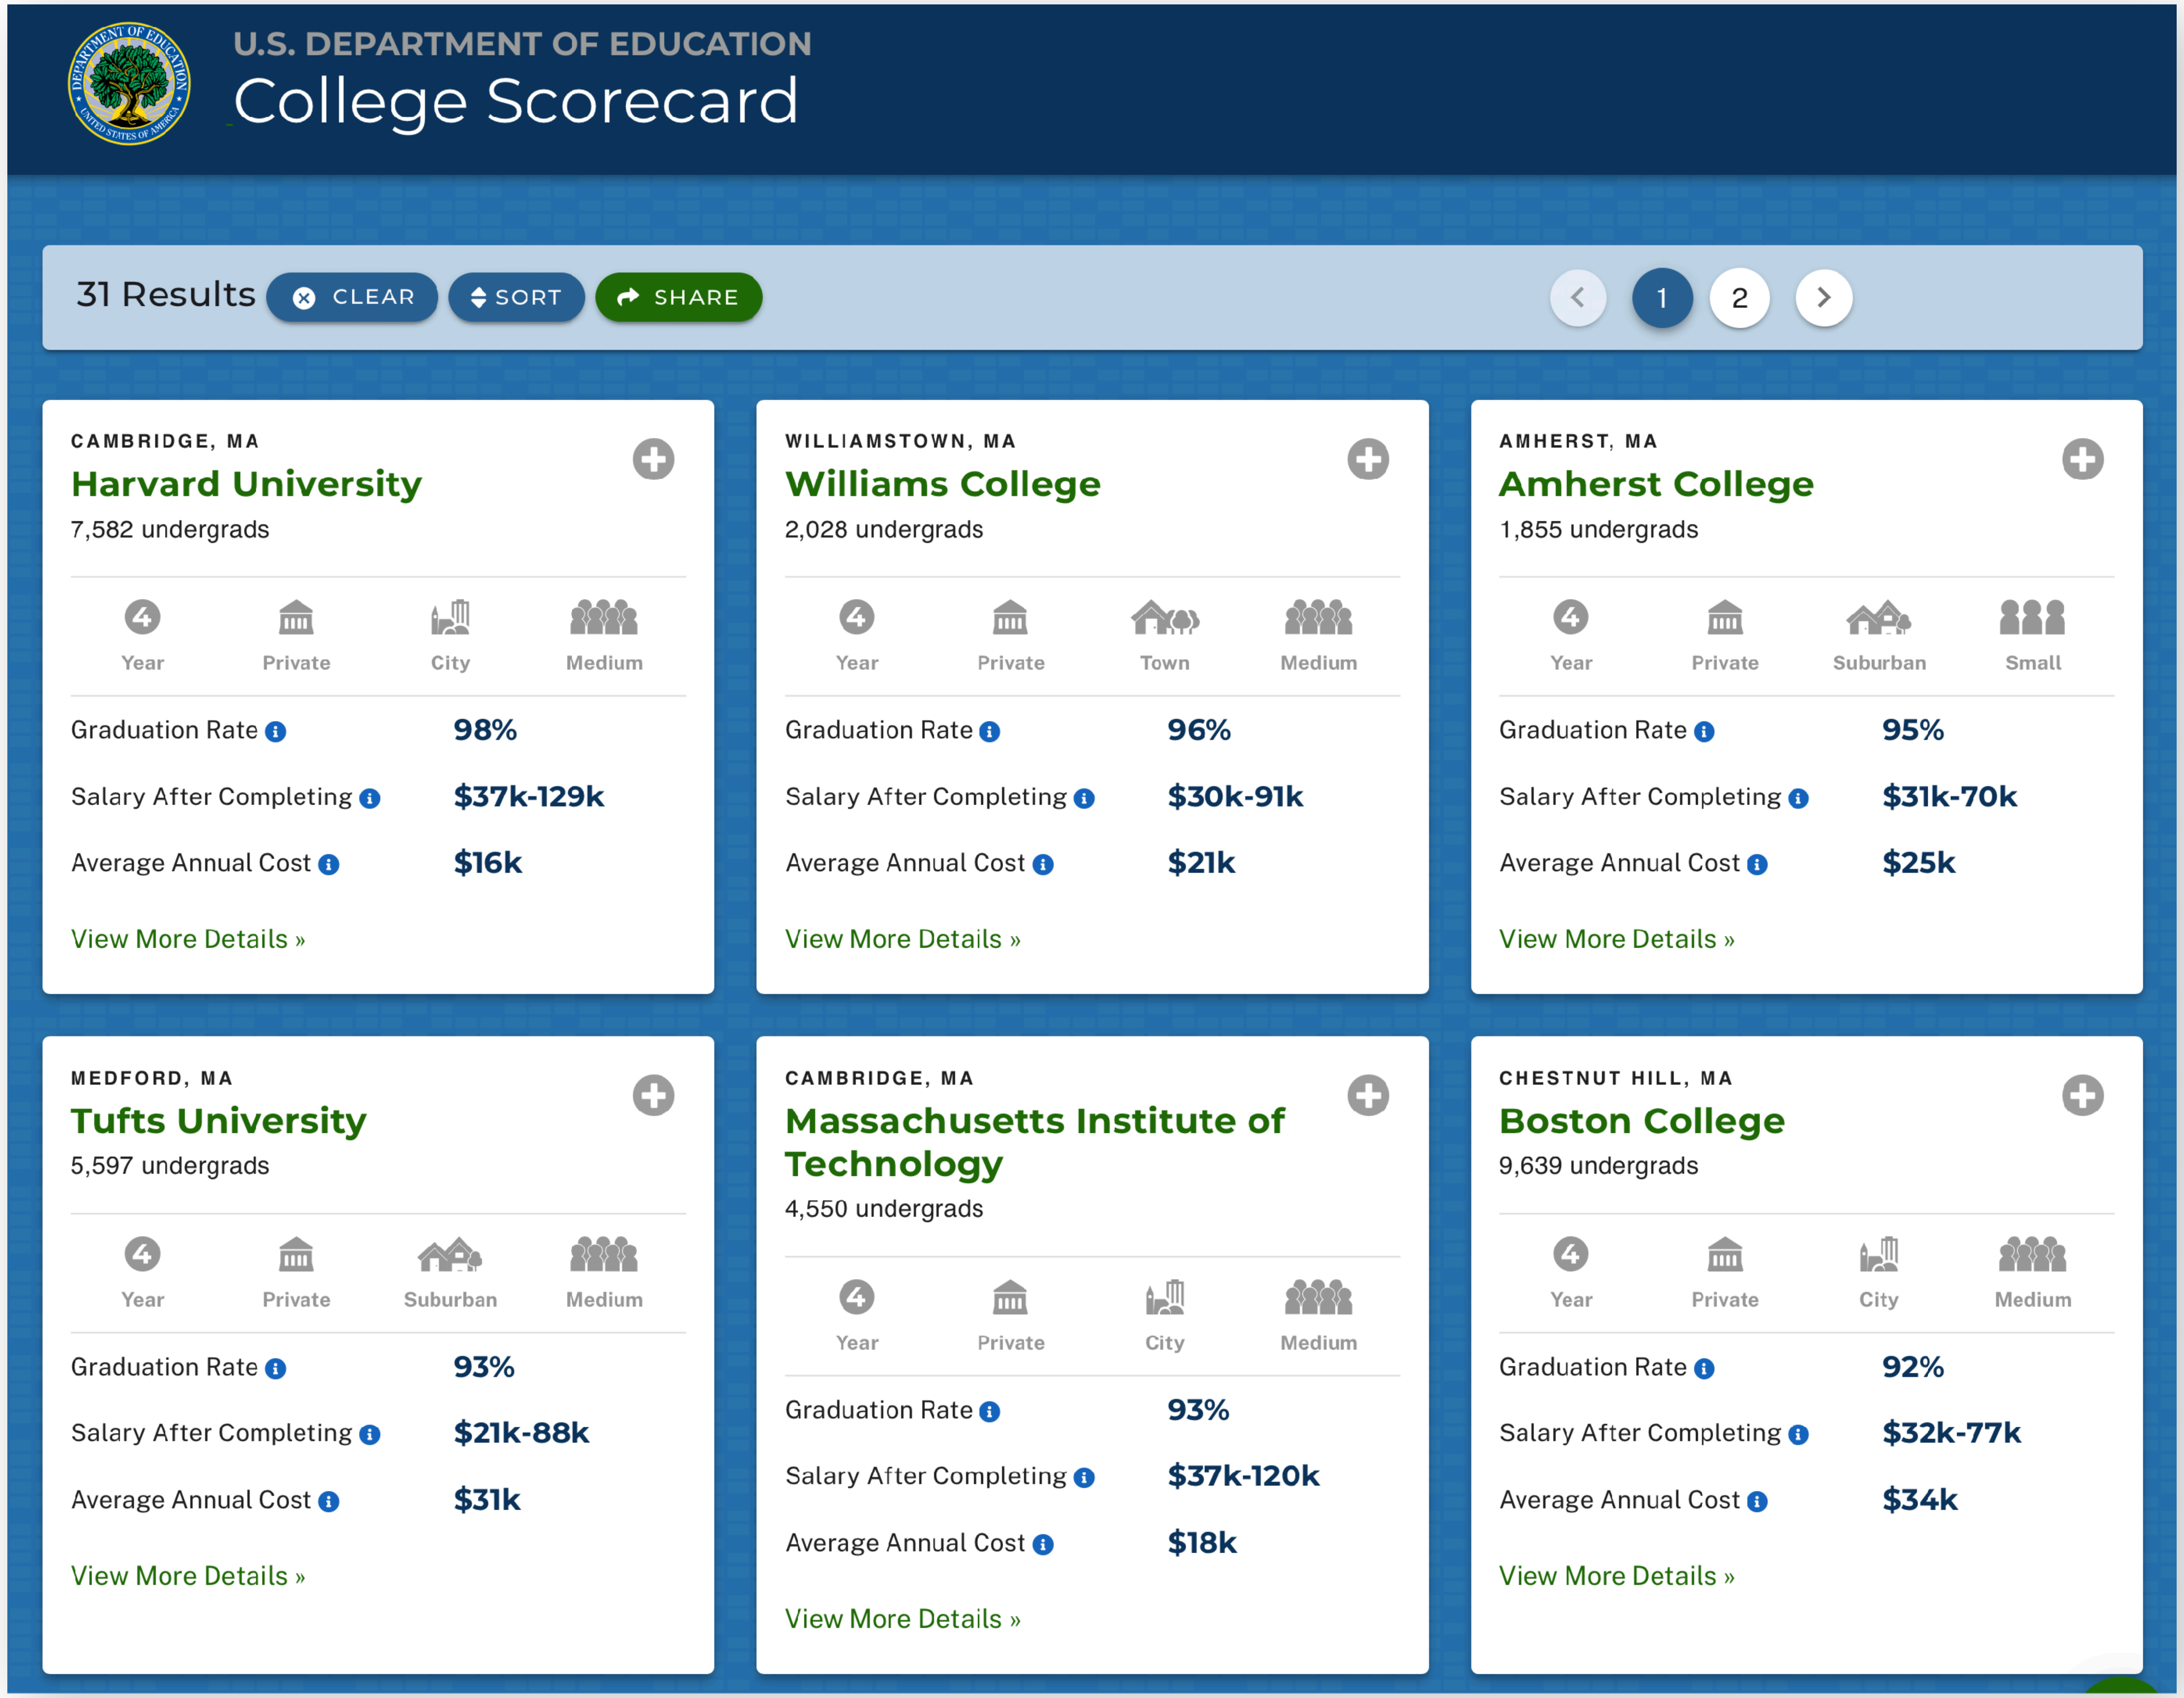

TUMULT LABS

Data Custodian

Internal Revenue Service

Has earnings data (~150m tax payers)

Data Analyst

Department of Education

Wants (a view of) the data

Request

College
DUKE UNIVERSITY

Degree
BACHELORS

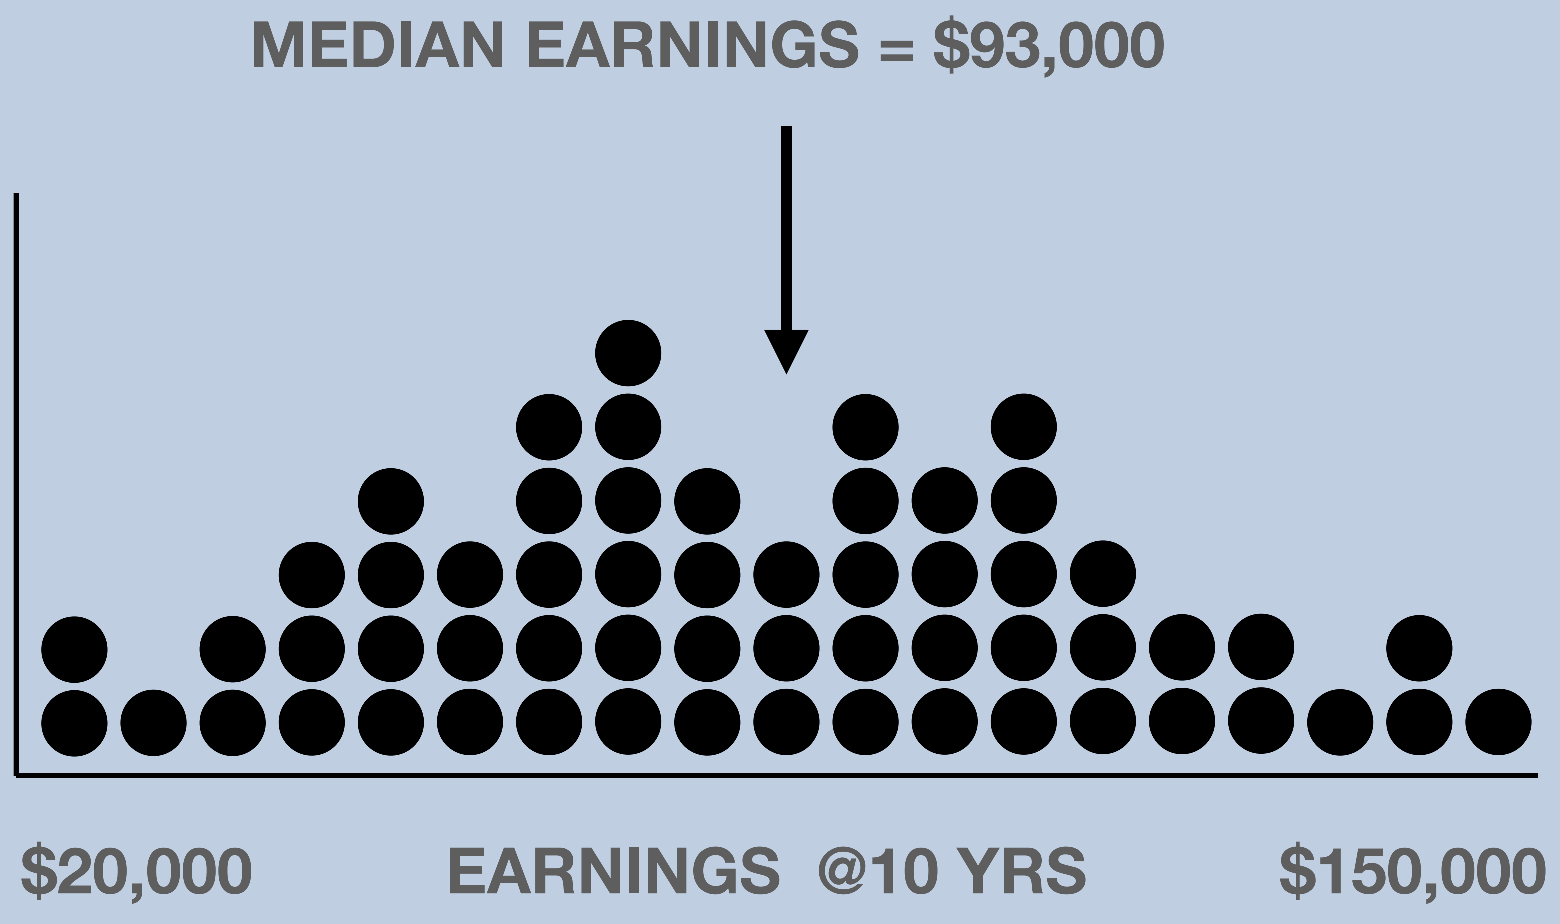

Response

College	Degree	Earnings (p50)
Duke	Bachelors	$93,000
Dartmouth	Bachelors	$92,000
Drexel	Bachelors	$76,000
…	…	…

College	Degree	Earnings (p50)
Duke	Bachelors	$93,000

Data Custodian

Internal Revenue Service

Must comply with regulation (US Title 26)
Bound by law to protect all information provided on tax returns (even fact of filing).

Must avoid privacy attacks

Data Analyst

Department of Education

Has defined and prioritized analytic tasks

Can describe “fitness-for-use” standards for tasks

Data Custodian

Internal Revenue Service

Data Analyst

Department of Education

Must comply with regulation
Bound by law to protect all information provided on tax returns (even fact of filing).

Must avoid privacy attacks

Has defined and prioritized analytic tasks

Can describe "fitness-for-use" standards for tasks

Lower Risk — Higher Risk

Lower data quality — Higher data quality

Bad outcome:
Lost insights, inability to complete analysis, incorrect conclusions, faulty decision-making

Bad outcome:
Privacy breach, violation of regulation, loss of institutional trust

TUMULT LABS

Informal privacy protection methods

“Informal” privacy protection:

(1) Ad-hoc distortion of income statistics

(2) Suppression of all statistics for groups deemed too small

2015

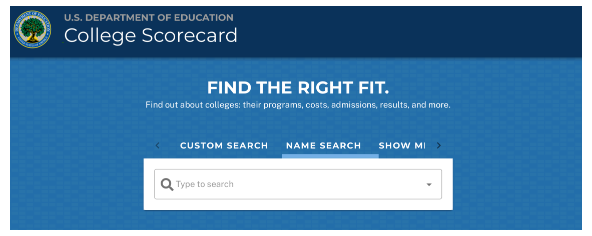

2016

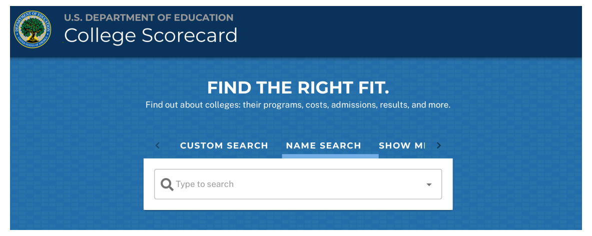

2017

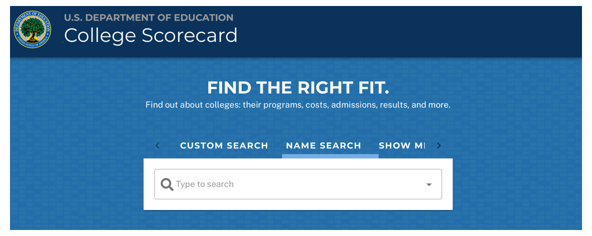

2018

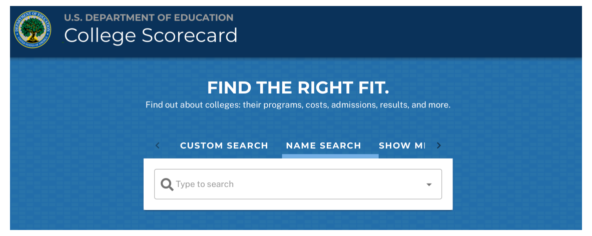

2019

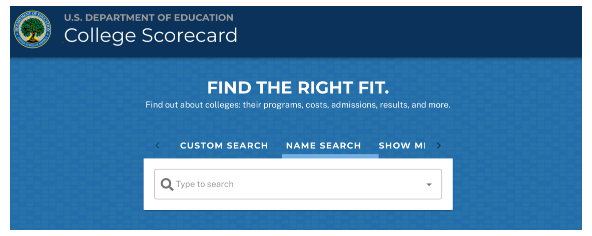

Adoption of differential privacy

"Informal" privacy protection:

(1) Ad-hoc distortion of income statistics

(2) Suppression of all statistics for groups deemed too small

Differential privacy

2015

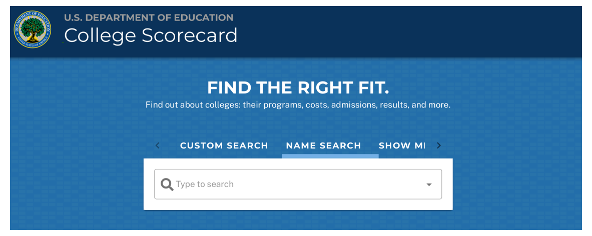

2016

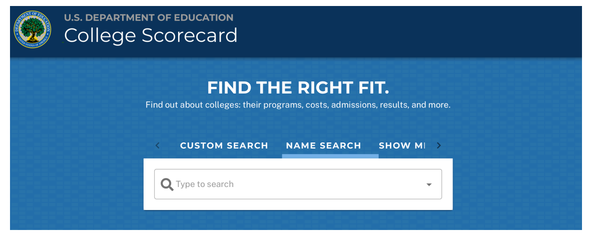

2017

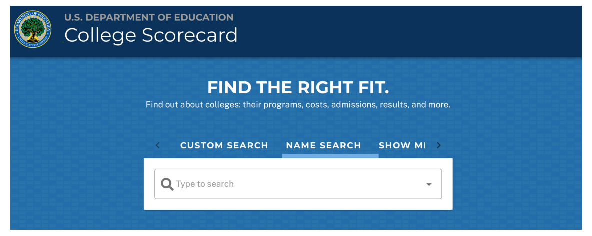

2018

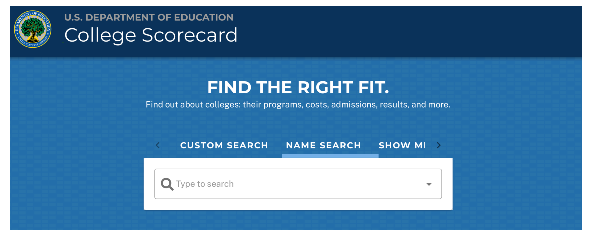

2019

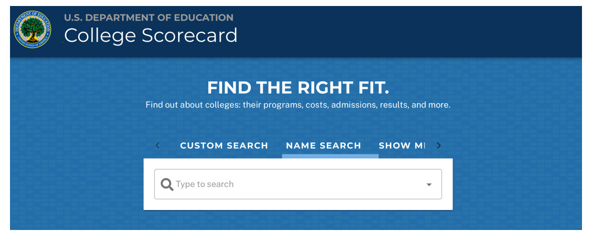

Feasibility Study

2020

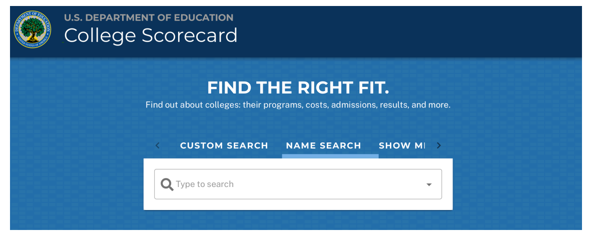

Published

2021

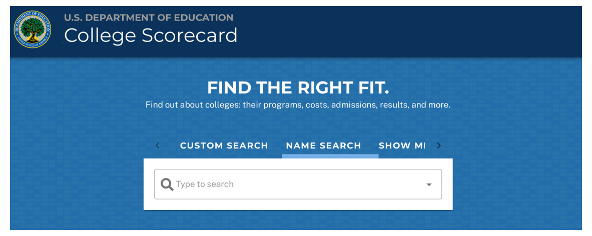

Published

2022

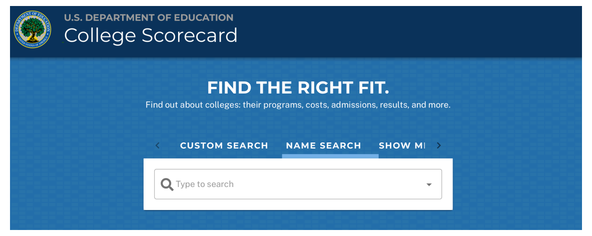

In Process

Steadily increasing requests for data

“Informal” privacy protection:

(1) Ad-hoc distortion of income statistics

(2) Suppression of all statistics for groups deemed too small

2015

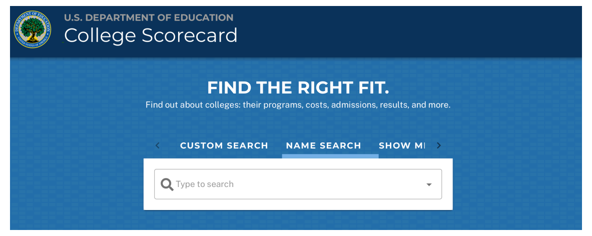

2016

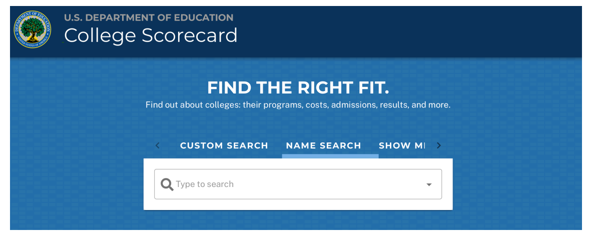

2017

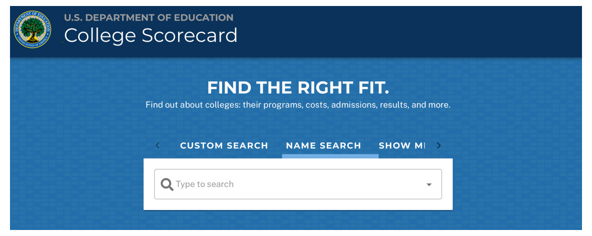

2018

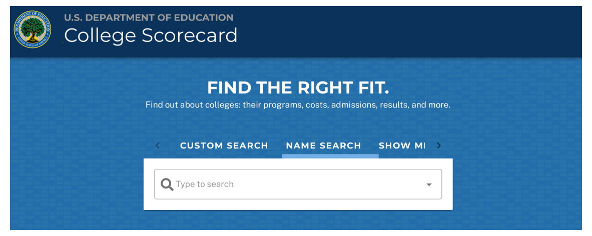

2019

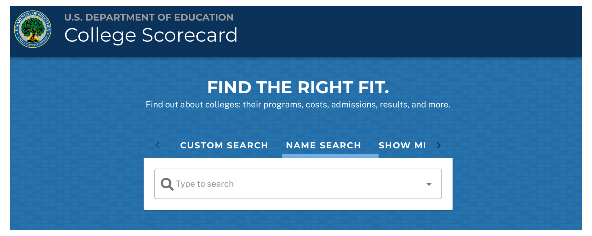

From INSTITUTION level To PROGRAM level

“Breakouts” by GENDER and PELL STATUS

From MEDIAN (P50) To P25, P50, and P75

COUNTS Students earning above 1.5 * Poverty Threshold

Increased risk for the data custodian

“Informal” privacy protection:

(1) Ad-hoc distortion of income statistics

(2) Suppression of all statistics for groups deemed too small

2015

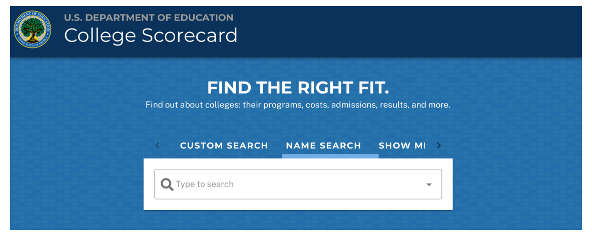

2016

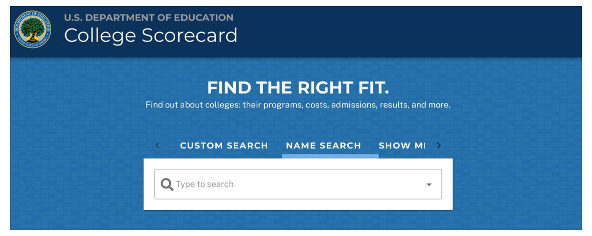

2017

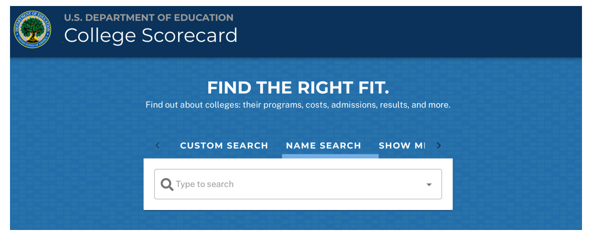

2018

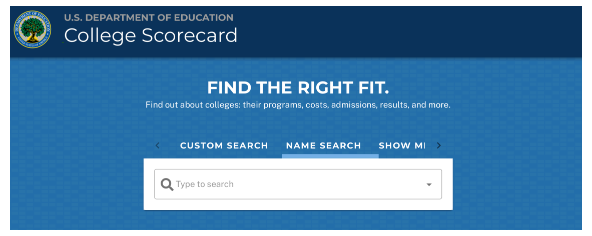

2019

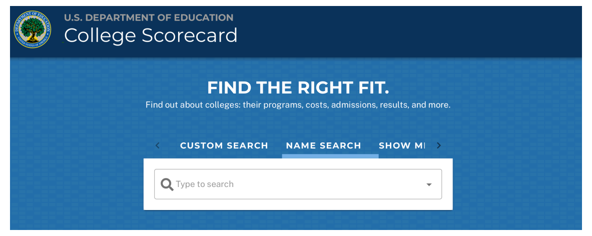

Tough questions for the data custodian

- **How much additional risk** for more detailed statistics?
- How much is my privacy risk growing with each annual release?
- What if one individual appears in multiple cohorts?
- **How should I respond**: how much more distortion? How much more suppression?

From INSTITUTION level To PROGRAM level

“Breakouts” by GENDER and PELL STATUS

From MEDIAN (P50) To P25, P50, and P75

COUNTS Students earning above 1.5 * Poverty Threshold

Adoption of differential privacy

"Informal" privacy protection:

(1) Ad-hoc distortion of income statistics

(2) Suppression of all statistics for groups deemed too small

Differential privacy

2015 2016 2017 2018 2019 2020 2021 2022

Differential privacy can help the custodian understand incremental risk and respond appropriately.

Differential privacy

a standard for computations on data
that limits the personal information that could be revealed by the output.

Guarantee of limited disclosure about input

FIRST	LAST	ZIP	SEX	AGE	ECOG	ICD-10
…	…	…	…	…	…	…
…	…	…	…	…	…	…
…	…	…	…	…	…	…
…	…	…	…	…	…	…
…	…	…	…	…	…	…
…	…	…	…	…	…	…
…	…	…	…	…	…	…
…	…	…	…	…	…	…
…	…	…	…	…	…	…

Sensitive individual-level data

DP analytics output

- Every individual protected.
- Every attribute protected.
- The guarantee holds, regardless of compute power or knowledge of potential attacker.
- Resists current and future attacks

A workflow for deploying differential privacy

Elicit Requirements

- Statistics requested
- Privacy/accuracy requirements

Prototype Algorithm

- Algorithmic strategy to compute released statistics with differential privacy

Identify Parameters

- Epsilon shares
- User contributions
- Suppression conditions

Interactively with analysts / stakeholders

- Adjust “levers” — algorithm parameters

- Visualize privacy loss vs fitness-for-use tradeoffs

Explore / Negotiate

Finalize & Deploy

- Finalize algorithm and parameters
- Deploy and generate final data product

Data release "levers"

Release description

☑ **Degree program**	P25	P50	P75
☑ **Pell=0 / Pell=1**	P25	P50	P75
☑ **Gender = 0 / Gender = 1**	P25	P50	P75

DIFFERENTIALLY PRIVATE ALGORITHM

Algorithm parameters

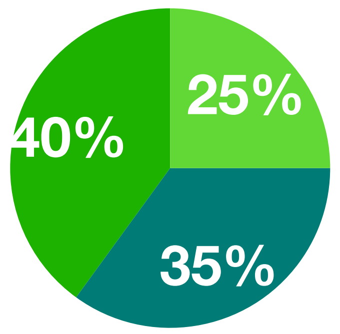

- ε_1
- ε_2
- ε_3

Suppression threshold: 15

Privacy semantics

Pure DP:
ε_total = 2.0

User contribution:
record

Source data

Tumult Platform

Outcome measures "fitness-for-use"

Measures describing "fitness-for-use"

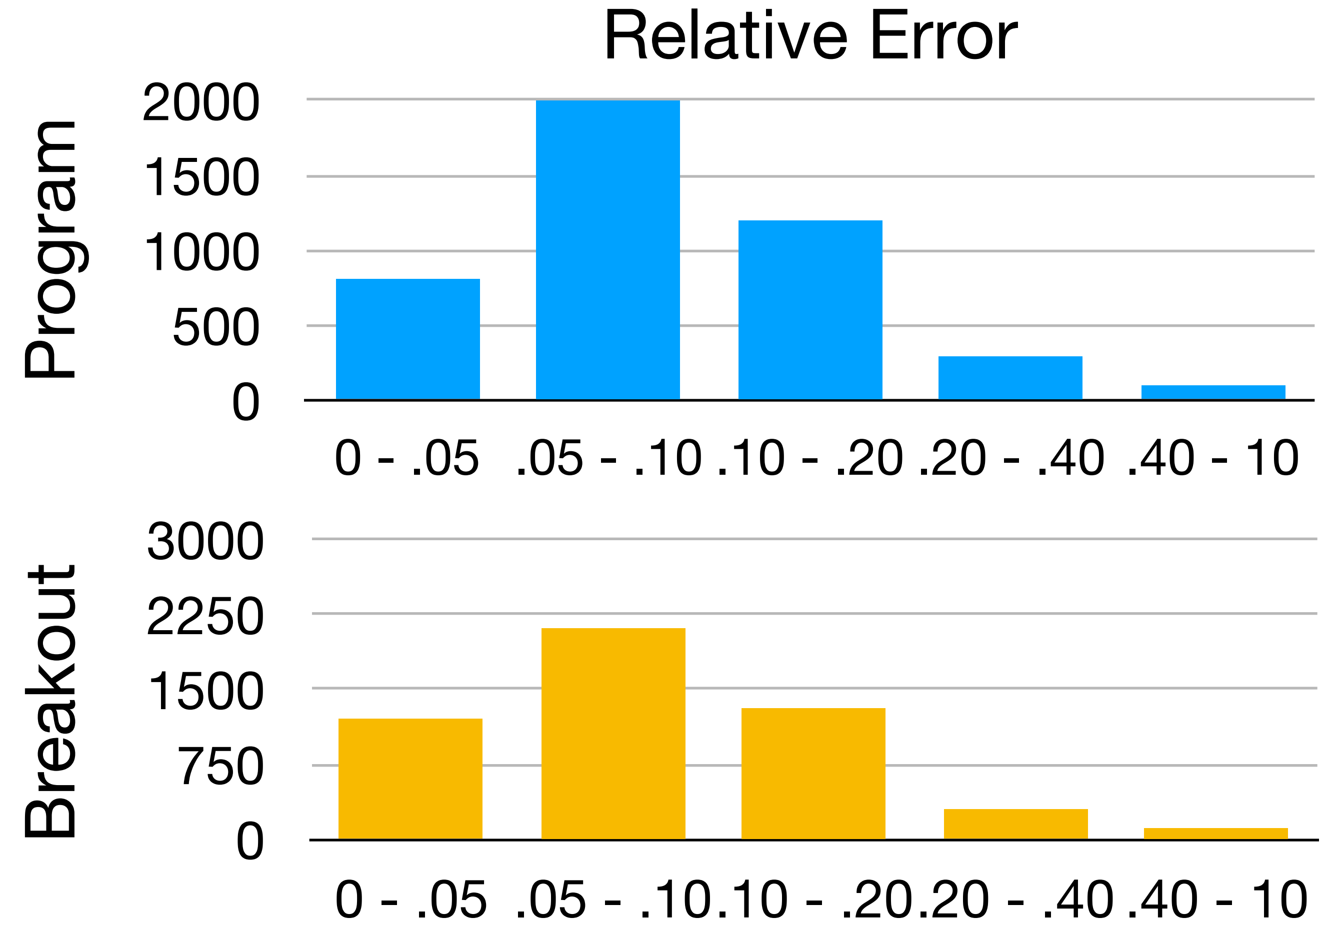

Suppressed groups
(% of total): 18

Suppressed groups
(% of total): 25

Data release "levers"

Release description

☑ **Degree program**	P25	P50	P75
☑ **Pell=0 / Pell=1**	P25	P50	P75
☑ **Gender = 0 / Gender = 1**	P25	P50	P75

DIFFERENTIALLY PRIVATE ALGORITHM

Algorithm parameters

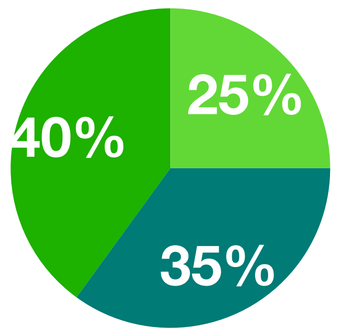

- ε_1
- ε_2
- ε_3

Suppression threshold: 15

Privacy semantics

Pure DP:
ε_total = 2.0

User contribution:
record

Source data

Tumult Platform

Outcome measures "fitness-for-use"

Measures describing "fitness-for-use"

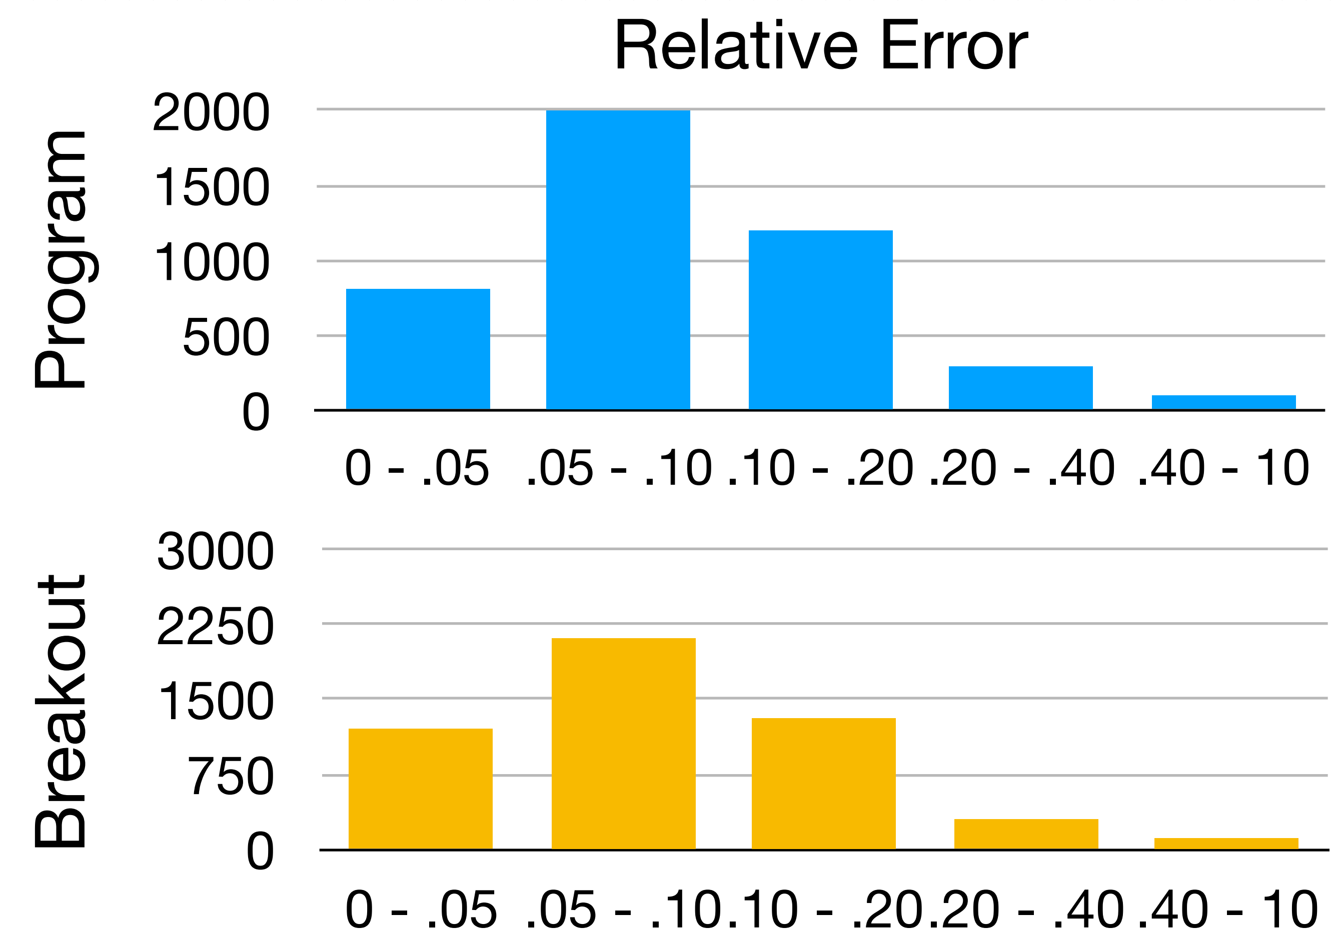

Suppressed groups
(% of total): 18

Suppressed groups
(% of total): 25

NO CONTROVERSY —Custodian & Analyst Both Win!

- **Are we using error-optimal DP algorithms?**
- **Can we get more data?**

Data release "levers"

Outcome measures "fitness-for-use"

Release description

☑ Degree program	P25	P50	P75
☑ Pell=0 / Pell=1	P25	P50	P75
☑ Gender = 0 / Gender = 1	P25	P50	P75

DIFFERENTIALLY PRIVATE ALGORITHM

Algorithm parameters

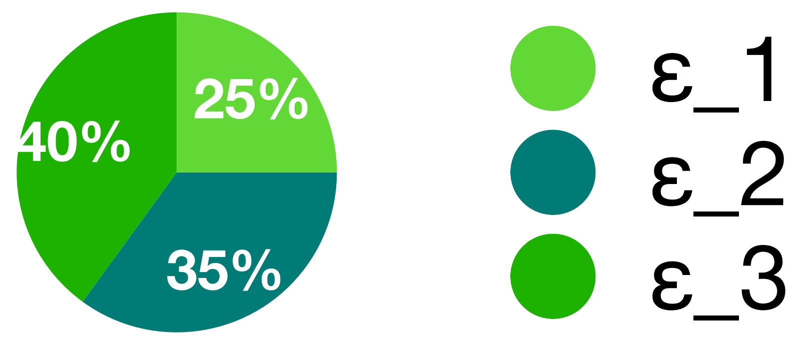

Suppression threshold: 15

Privacy semantics

Pure DP:
ε_total = 2.0

User contribution:
record

Source data

Tumult Platform

Measures describing "fitness-for-use"

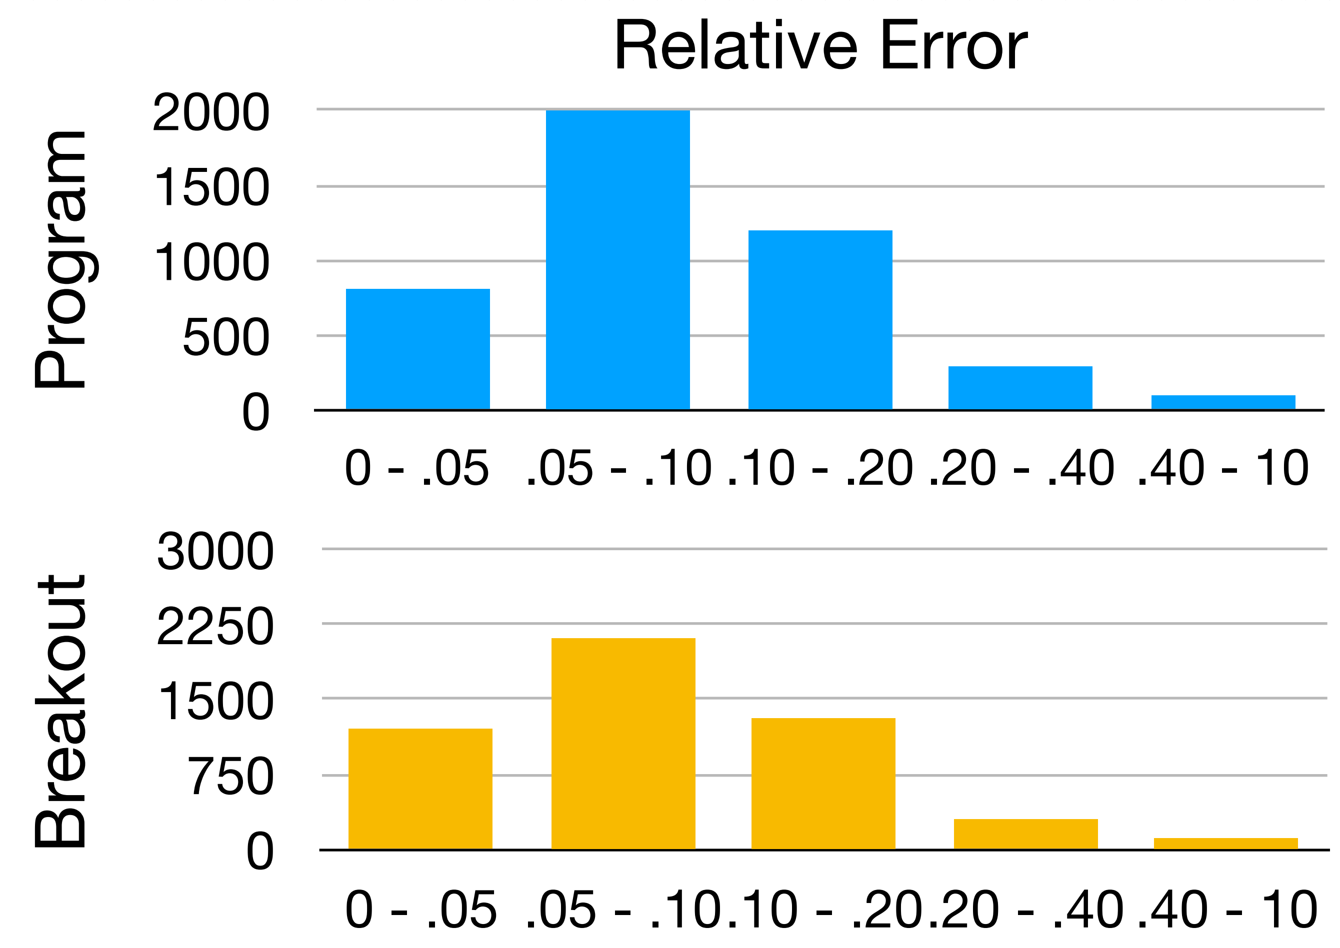

Suppressed groups
(% of total): 18

Suppressed groups
(% of total): 25

Analyst can add or remove to the released statistics

Custodian sets bound on privacy loss

Analyst can adjust algorithm parameters

Data Custodian
Internal Revenue Service

Data Analyst
Department of Education

Outcomes

Utility
More student earnings statistics than previous releases, with comparable accuracy.

Assurance and risk management
A rigorous, quantifiable privacy guarantee to guide decision-making about privacy risk.

Ease-of-use
Streamlined communication about privacy / accuracy tradeoffs.

Conclusions and challenges

- Differential privacy encourages custodians and analysts to carefully consider data uses and fitness-for-use standards.
  - A move from “universal” data products to customized data products.
- Tools to support iterative exploration and negotiation are essential, but don’t exist in most privacy platforms.
- Calculating and communicating error to analysts and stakeholders is challenging (and could incur its own privacy loss!)
- Data consumers don’t want to see high error outputs; they prefer them to be suppressed, even when error is quantified.

Thank you!

Questions?

www.tmlt.io/connect

miklau@tmlt.io