Exploring the Limits of Differential Privacy

TPRC52

September 20, 2024

Washington, DC

David D. Clark (MIT CSAIL)
Simson Garfinkel (Harvard John A. Paulson School of Engineering and Applied Sciences)
KC Claffy (University of California, San Diego)

Clark, David D. and Garfinkel, Simson and Claffy, KC C., Exploring the Limits of Differential Privacy (July 31, 2024). Available at SSRN: https://ssrn.com/abstract=4911177

Differential privacy — it's the future.

Invented in 2006 and used in the US 2020 Census.

Widely recognized as useful and powerful privacy-enhancing technology (PET).

Called for in "National strategy to advance privacy-preserving data sharing and analytics," NCO NITRD, Washington, DC, USA, Tech. Rep., Mar. 2023.

Provides *mathematical certainty* regarding maximum "privacy loss" for any data release.

Composable — Differential privacy avoids the "mosaic problem" that befuddles other privacy technologies like de-identification.

Tunable — Data curator can control the privacy loss/utility trade-off.

Worst Case Assumption — Protects outliers and everybody else.

Some funding agencies are encouraging researchers to use DP to release their data.

DP’s goal is to prevent database reconstruction

Differential privacy protects confidential data used for public statistics.

Example:

- You are in a class with 9 other students.
- The teacher announces that the average score is 98%.
- You look at your test and you got an 80%.

- Now you know the grades for everyone in the class…

ChatGPT

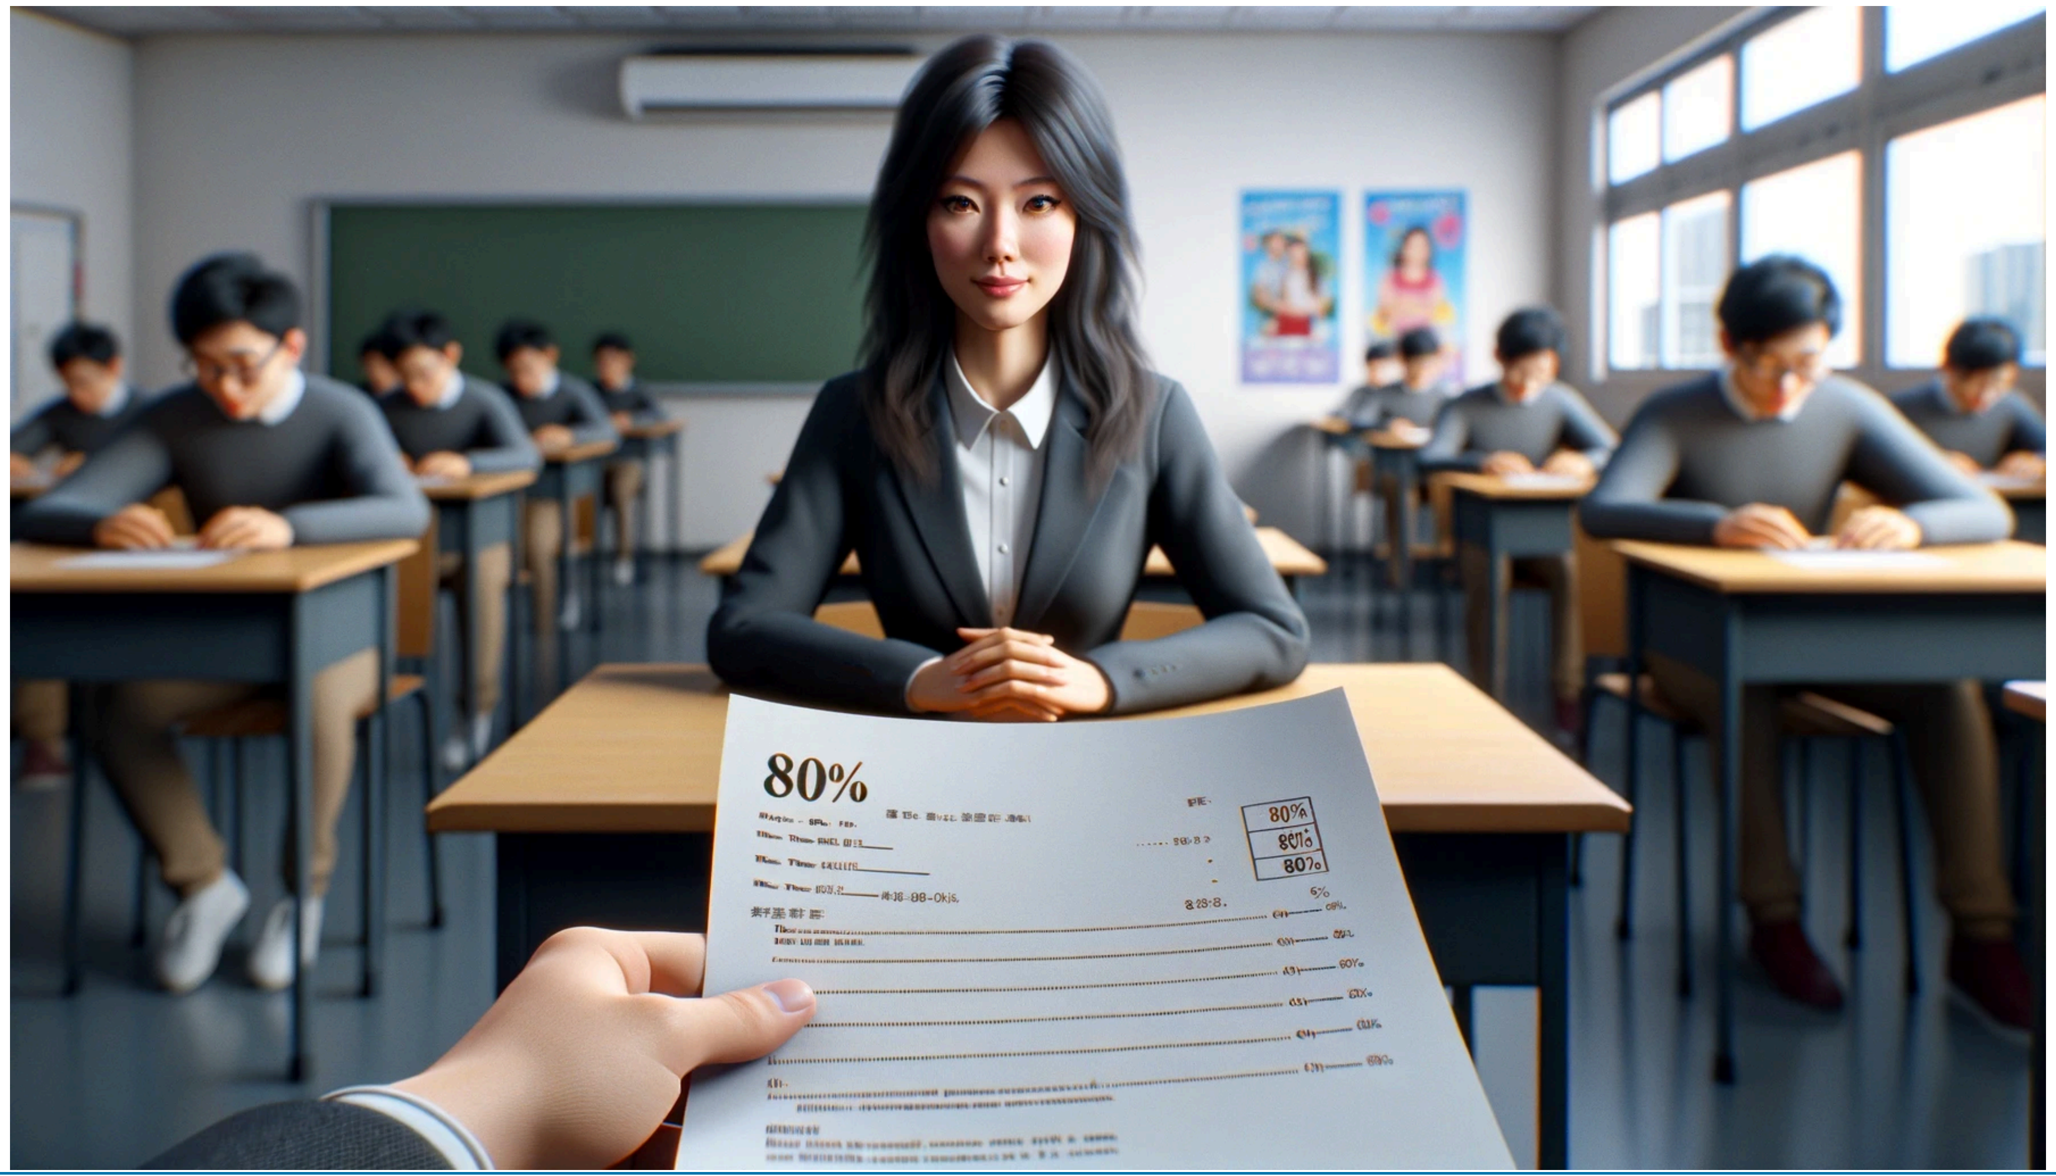

ChatGPT

Consider a survey of companies — what % of your systems are patched?

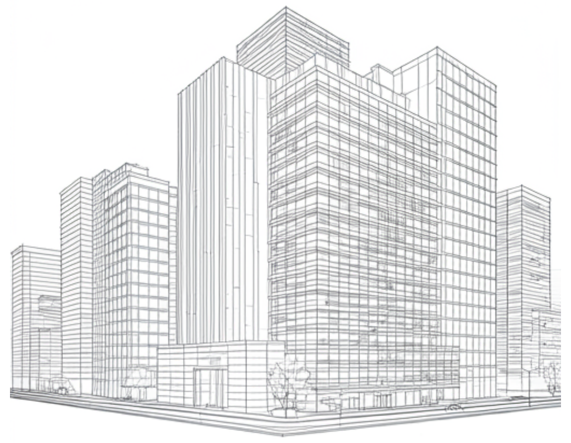

Accurate Statistics…

January

Company	# systems	% patched
Alpha	100	50
Bobble	100	50
Cantana	100	50
Delmax	100	50

Statistical Tabulation

Companies: 4
Average % patched: 50%

February

Company	# systems	% patched
Alpha	100	50
Bobble	100	50
Cantana	100	50
Delmax	100	50
Echo	100	25

Statistical Tabulation

Companies: 5
Average % patched: 45%

It's pretty easy to figure out that Echo has 25% of its systems patched

DP solves this problem by adding noise to published results

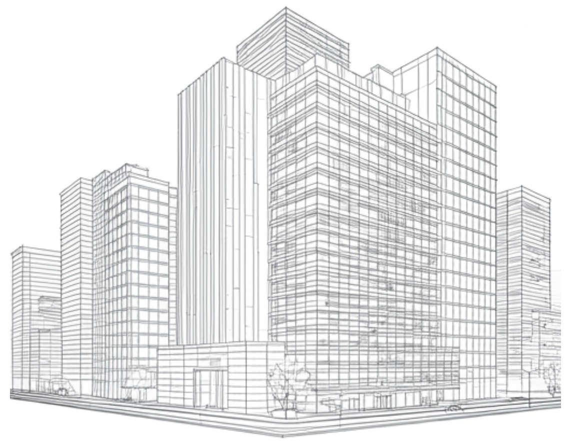

With Noise!

January

Company	# systems	% patched
Alpha	100	50
Bobble	100	50
Cantana	100	50
Delmax	100	50

Statistical Tabulation

Companies: 4
Average % patched: 55%

February

Company	# systems	% patched
Alpha	100	50
Bobble	100	50
Cantana	100	50
Delmax	100	50
Echo	100	25

Statistical Tabulation

Companies: 5
Average % patched: 47%

We don’t know what noise was added, so we can’t figure out Echo’s contribution.

How much noise is enough?

DP “Laplace Mechanism”

$$f^*(x) = f(x) + \mathrm{Lap}(\frac{\Delta f}{\epsilon})$$

f = function to make private

Lap = Laplace Noise

Δf = Sensitivity (how much each person can change the function)

ε = The privacy loss parameter. (0 = full privacy; ∞ = infinite privacy loss)

How much noise do we add? That's a policy decision.

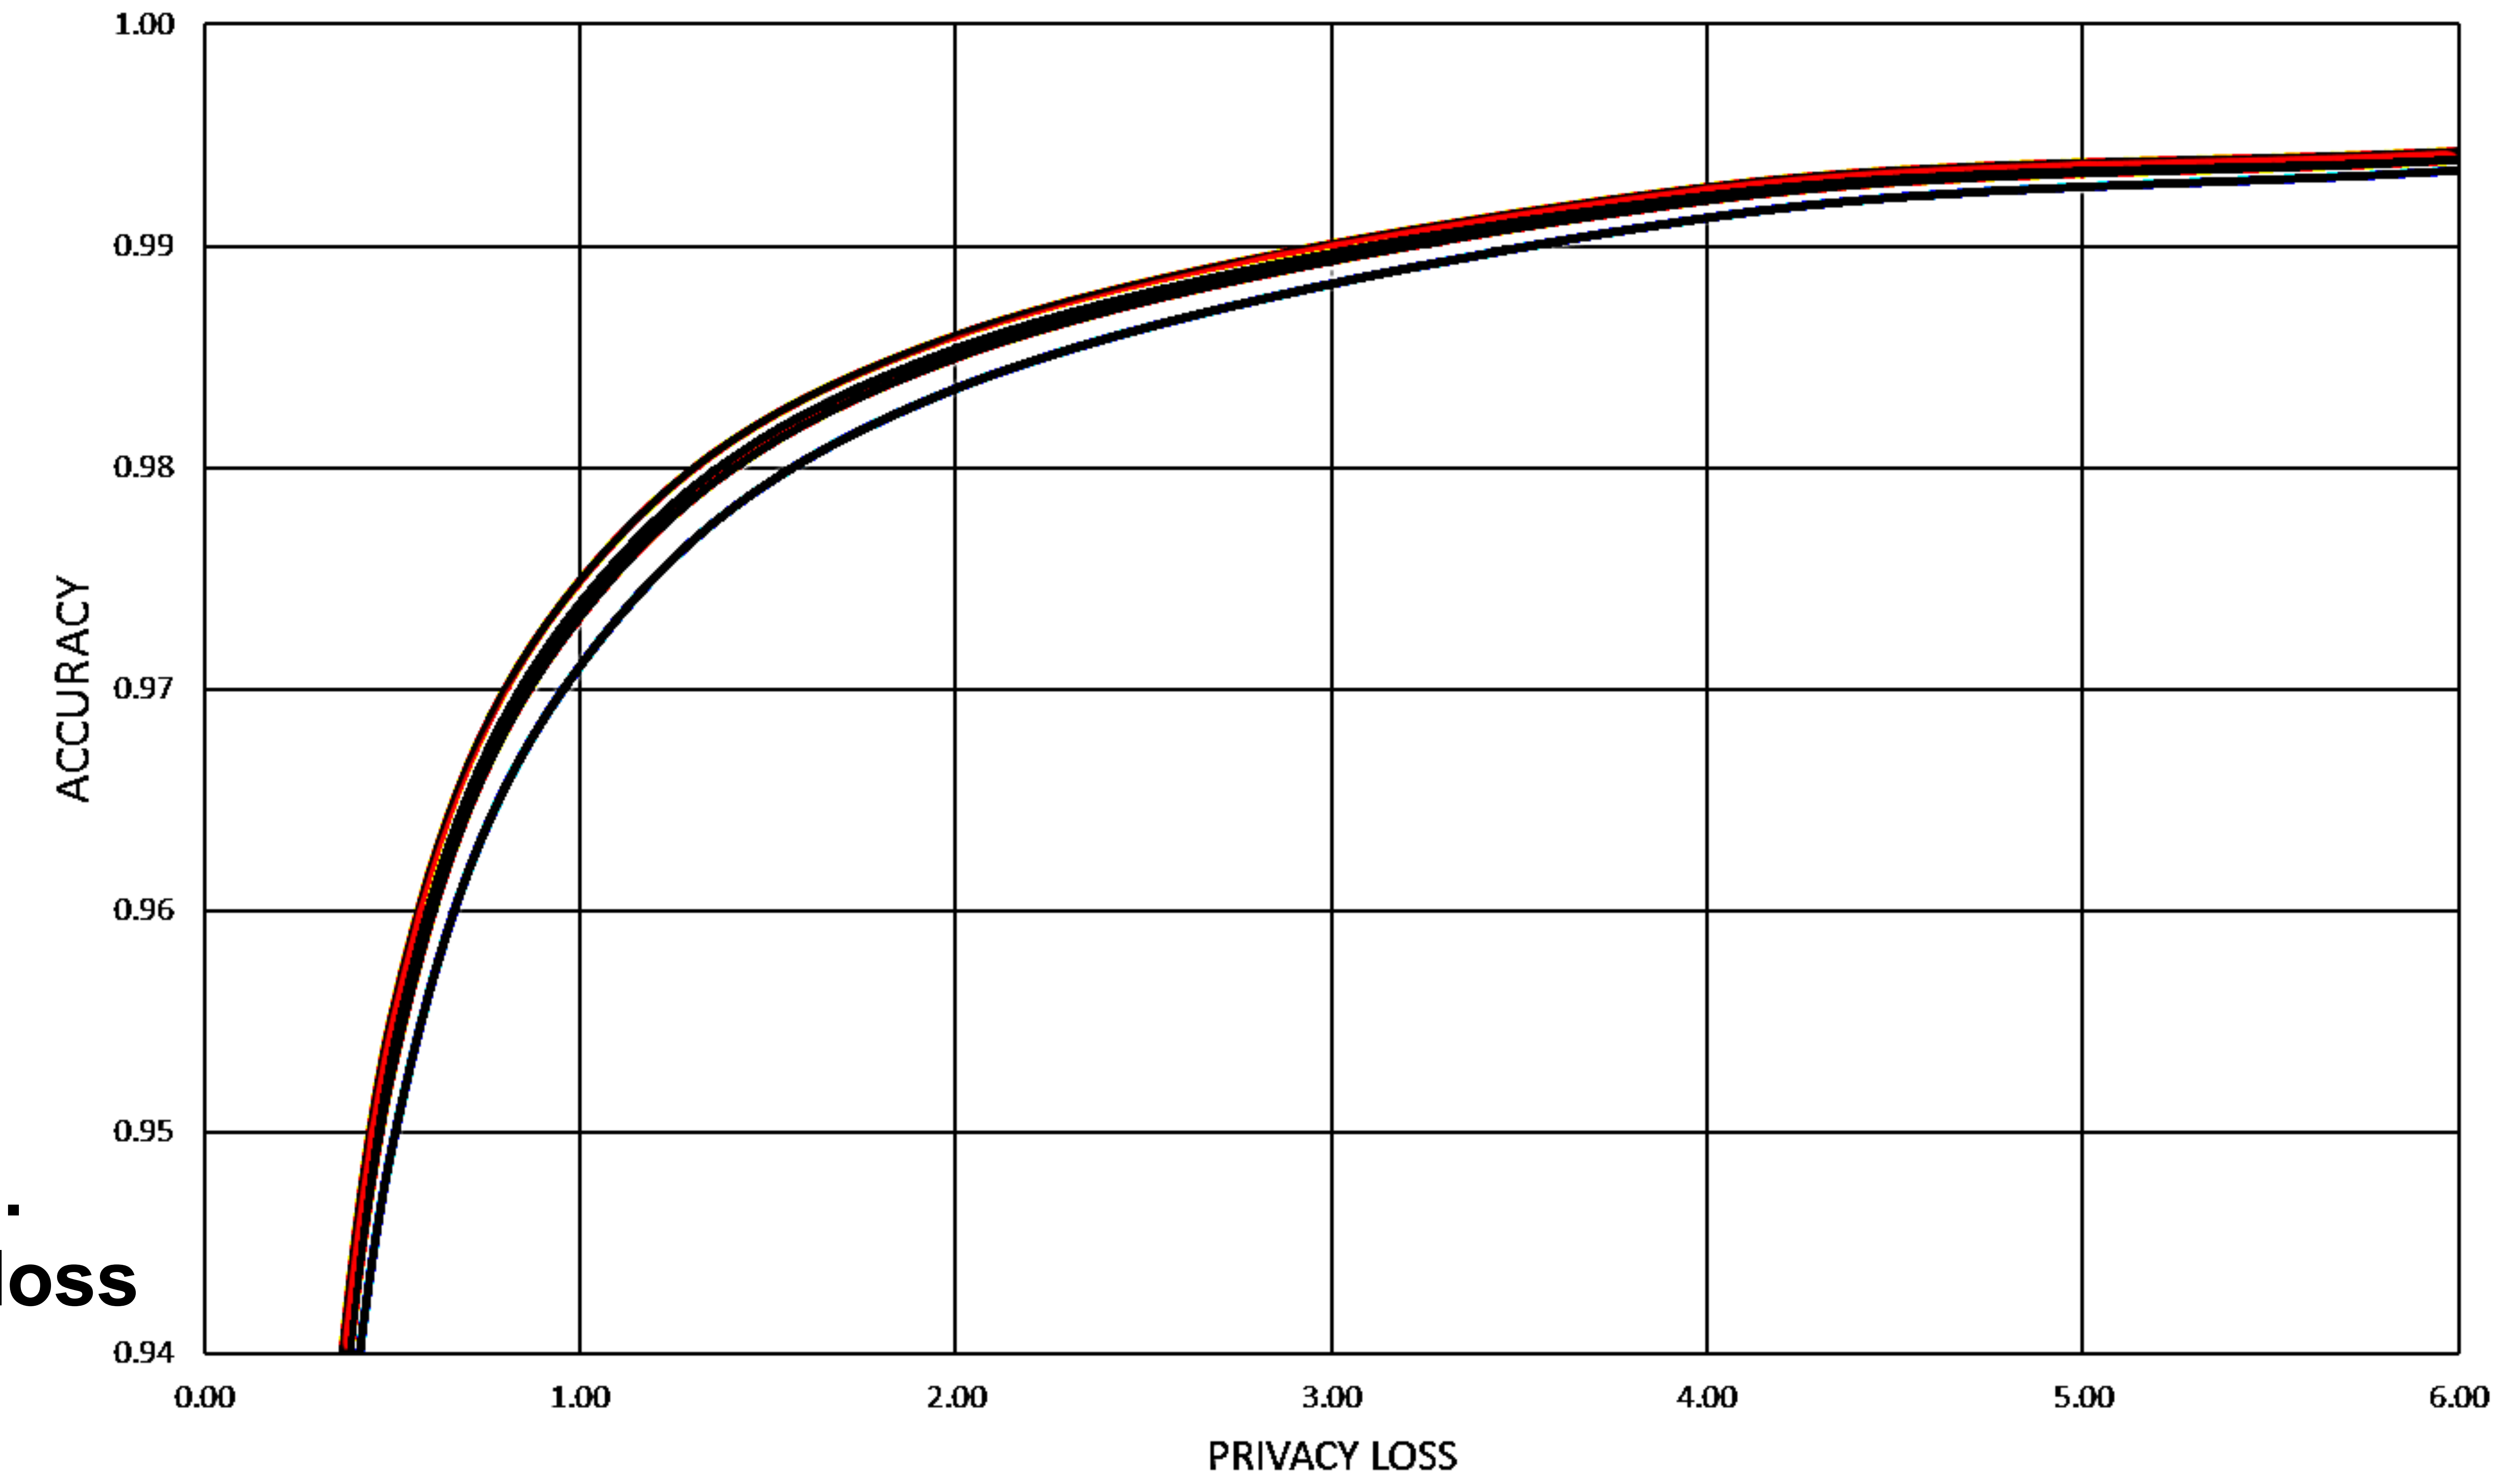

Highly accurate.
High privacy loss

Less accurate.
Little privacy loss

NIST SP 800-226 ipd
December 2023

Another way to look at the privacy loss/accuracy trade-off

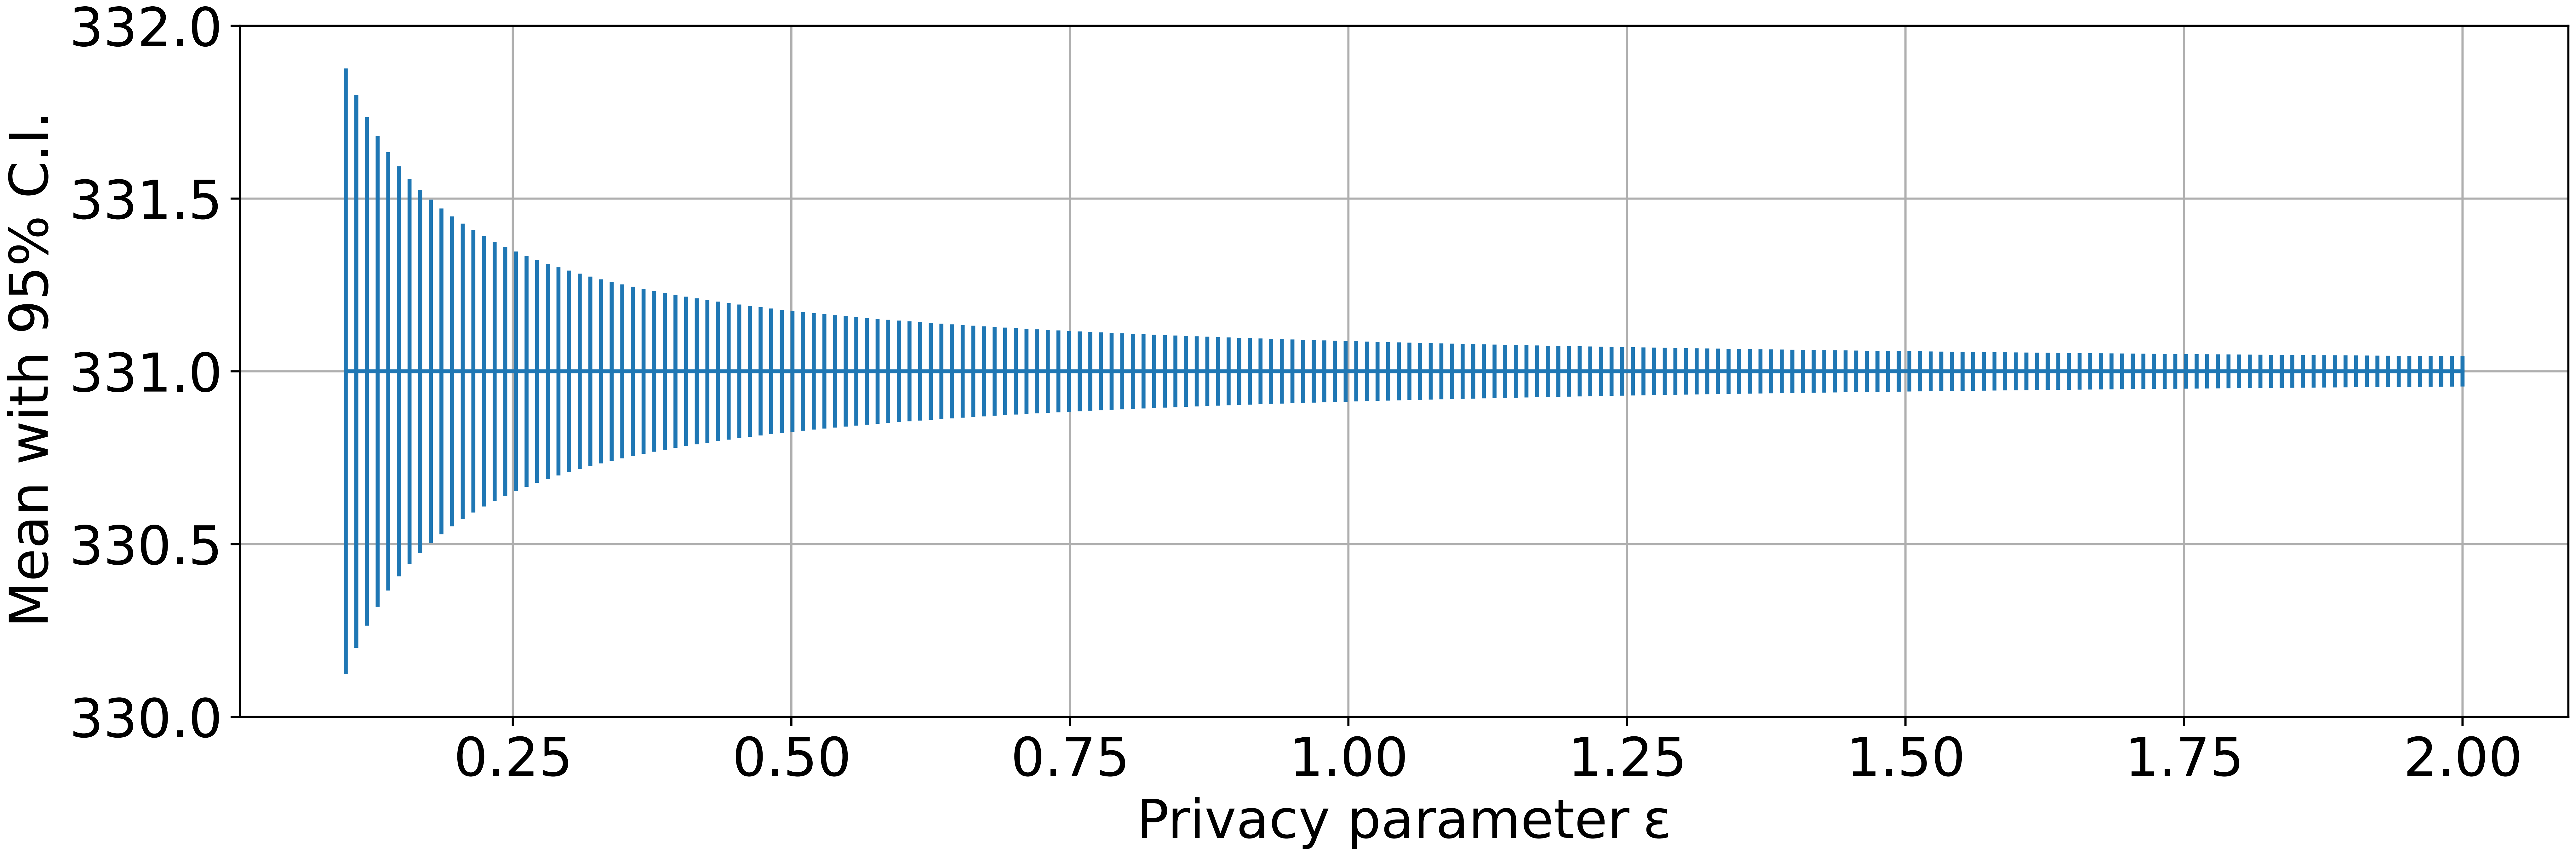

Fig. 7. The 95% confidence interval for the absolute error of the Laplace mechanism.

Ways of using DP — three models

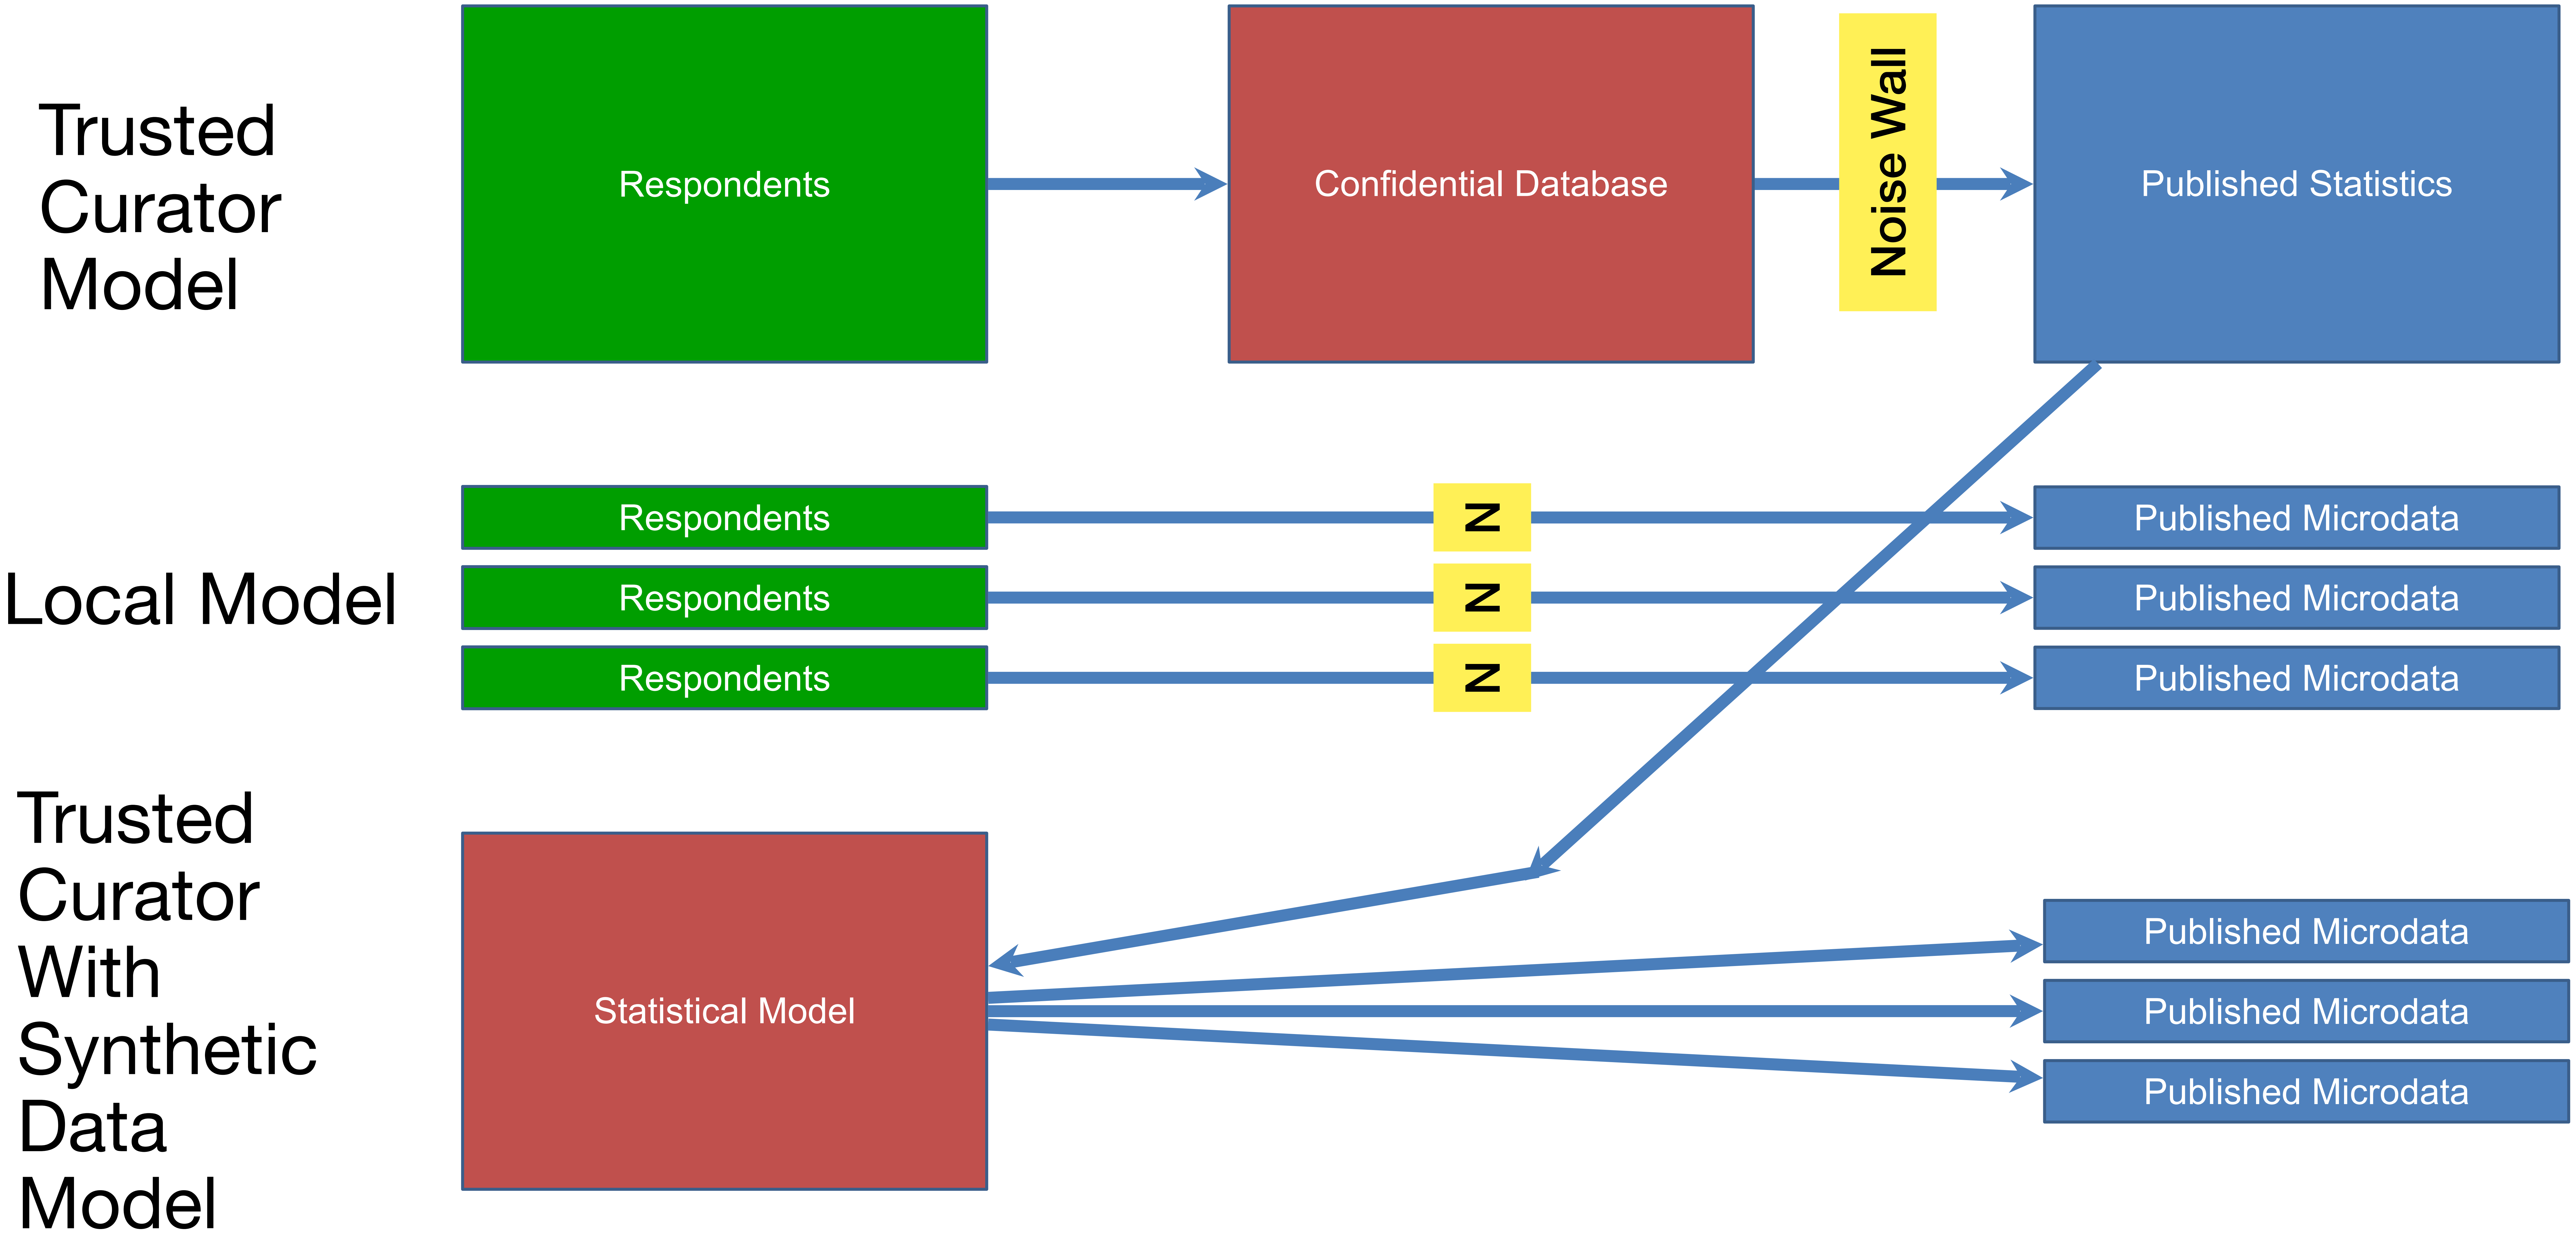

These examples use $\varepsilon=1$

Note $\varepsilon=1$ is almost always the wrong choice.

Pure DP uses Laplace noise.
What if there are 100 companies and they all have 50% patched?

Company	# systems	% patched
Alpha	100	50
Bobble	100	50
Cantana	100	50
Delmax	100	50
…		
Company 49	100	50
Company 50	100	50
…		
Company 100	100	100

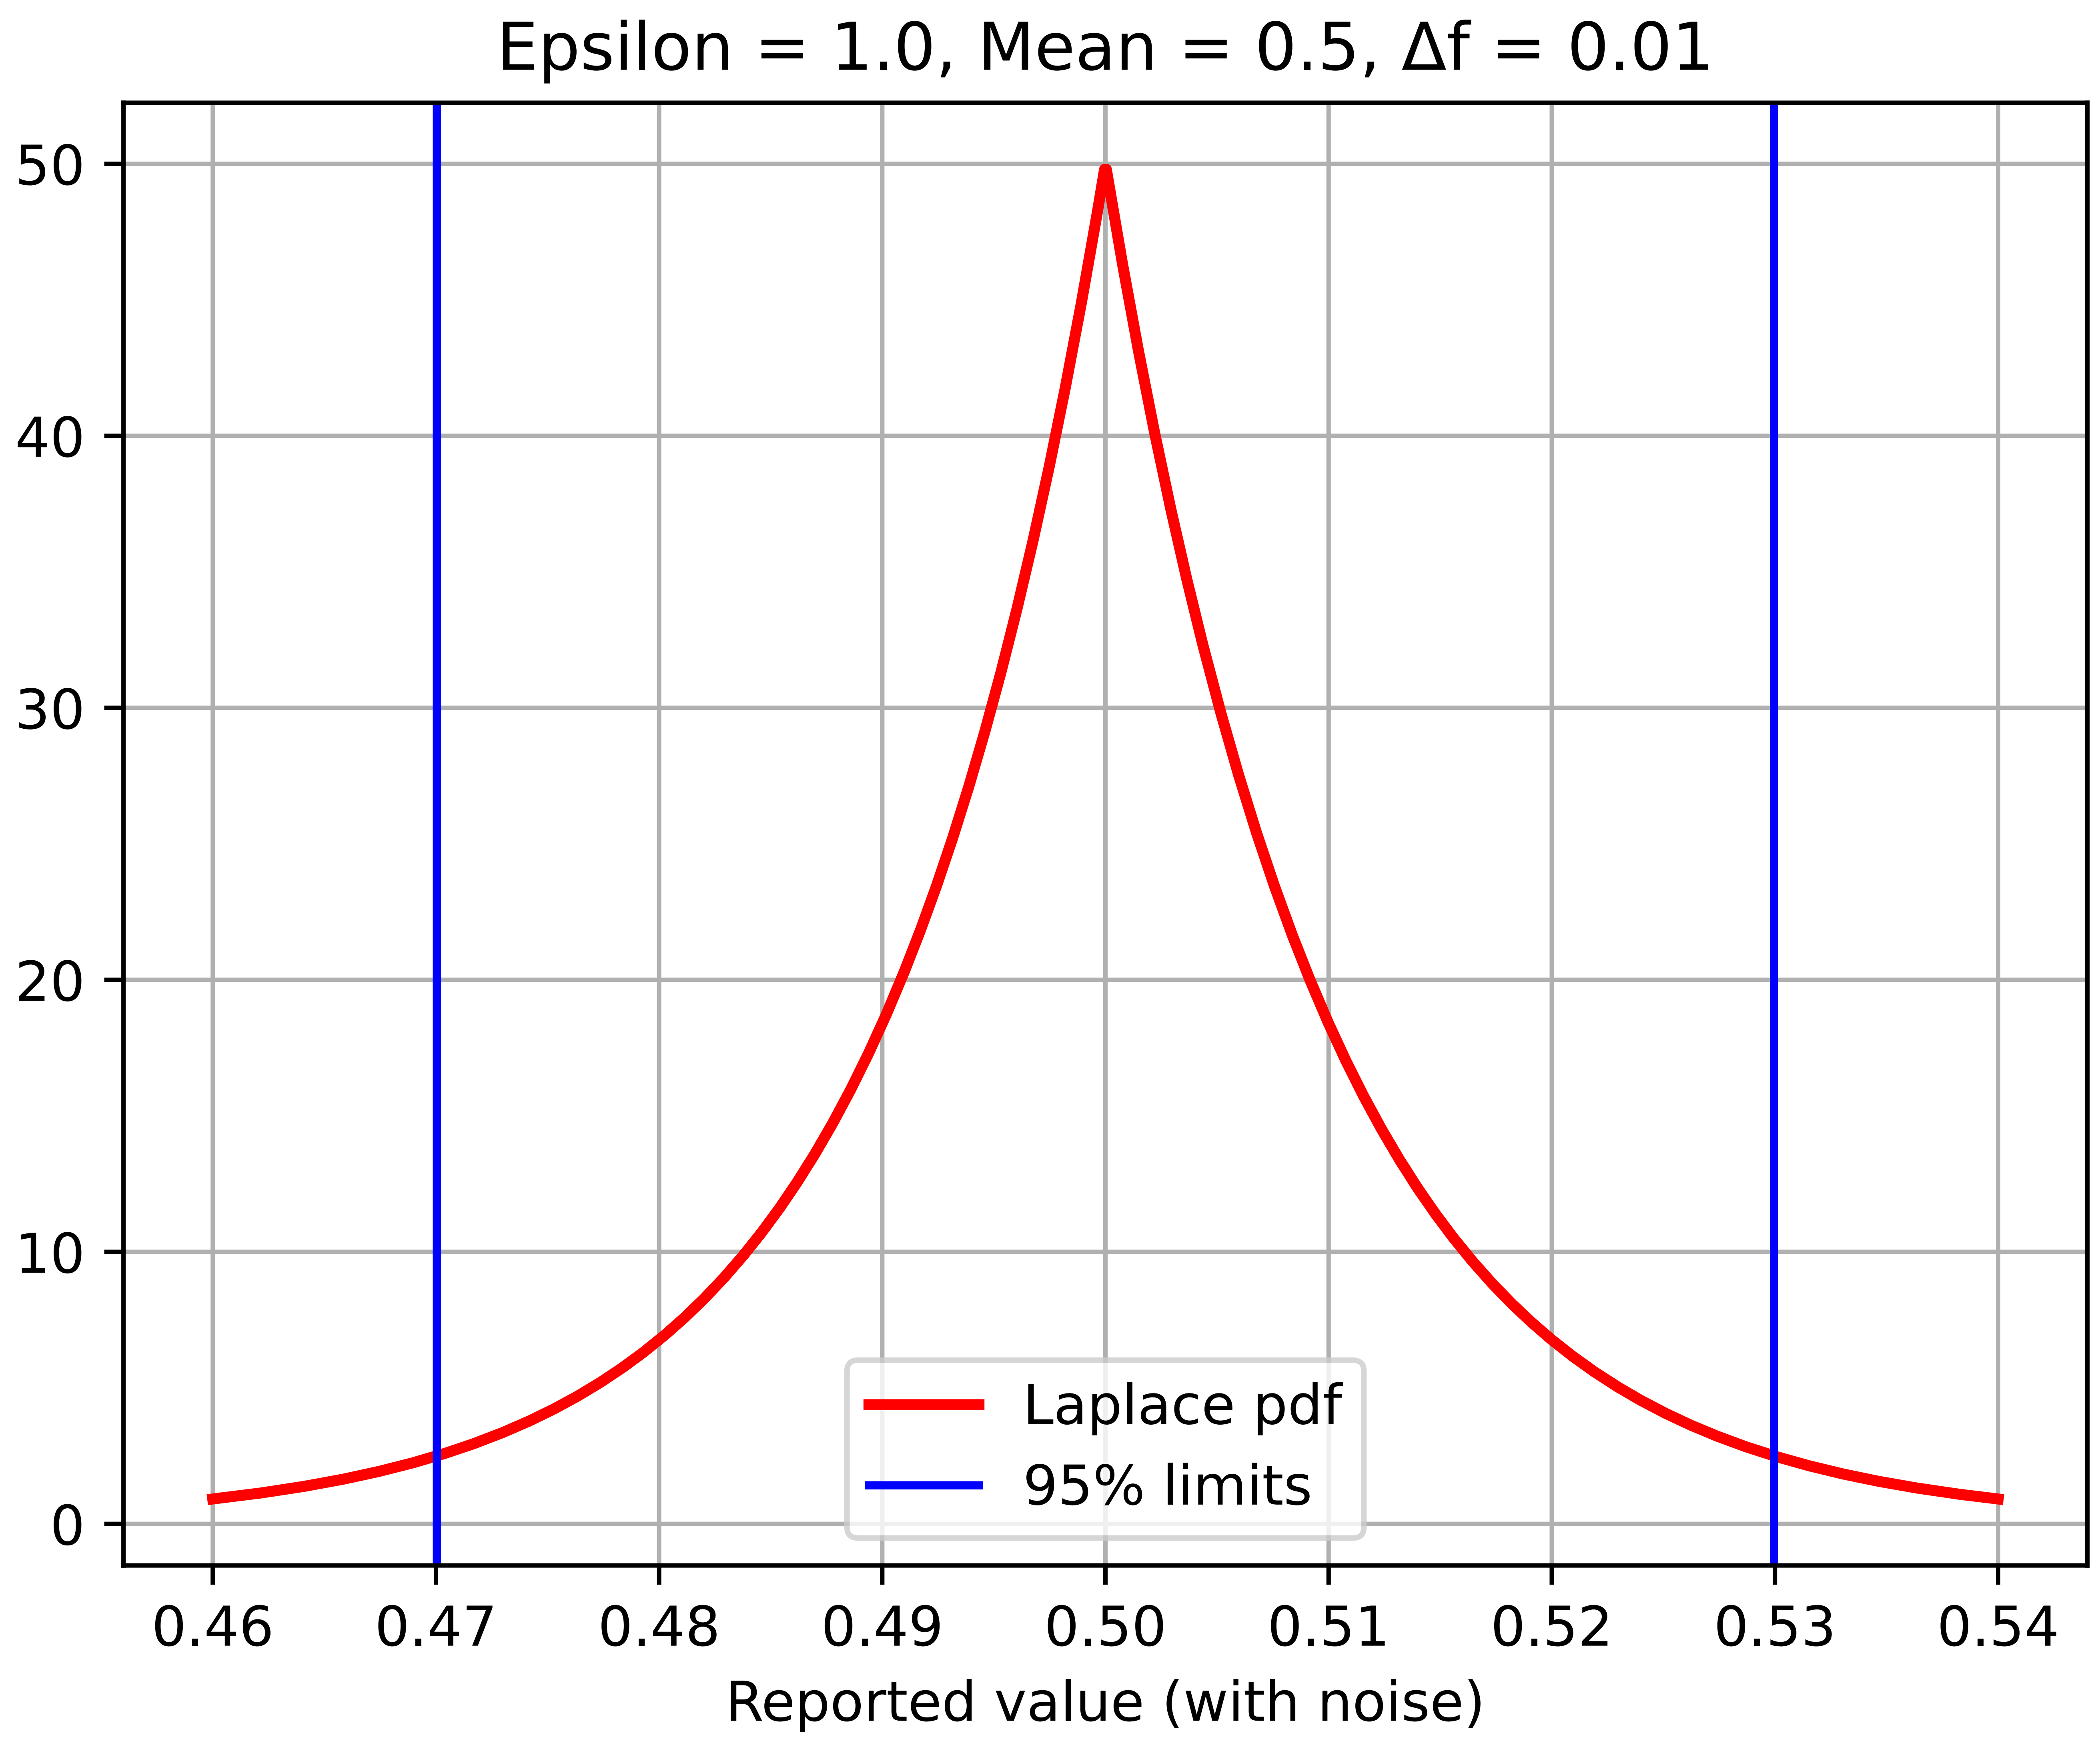

Note we are looking at just ε=1

It looks the same if there are 50 companies with 0% patched and 50 companies with 100% patched.

Company	# systems	% patched
Alpha	100	0
Bobble	100	0
Cantana	100	0
Delmax	100	0
…		
Company 49	100	0
Company 50	100	100
…		
Company 100	100	100

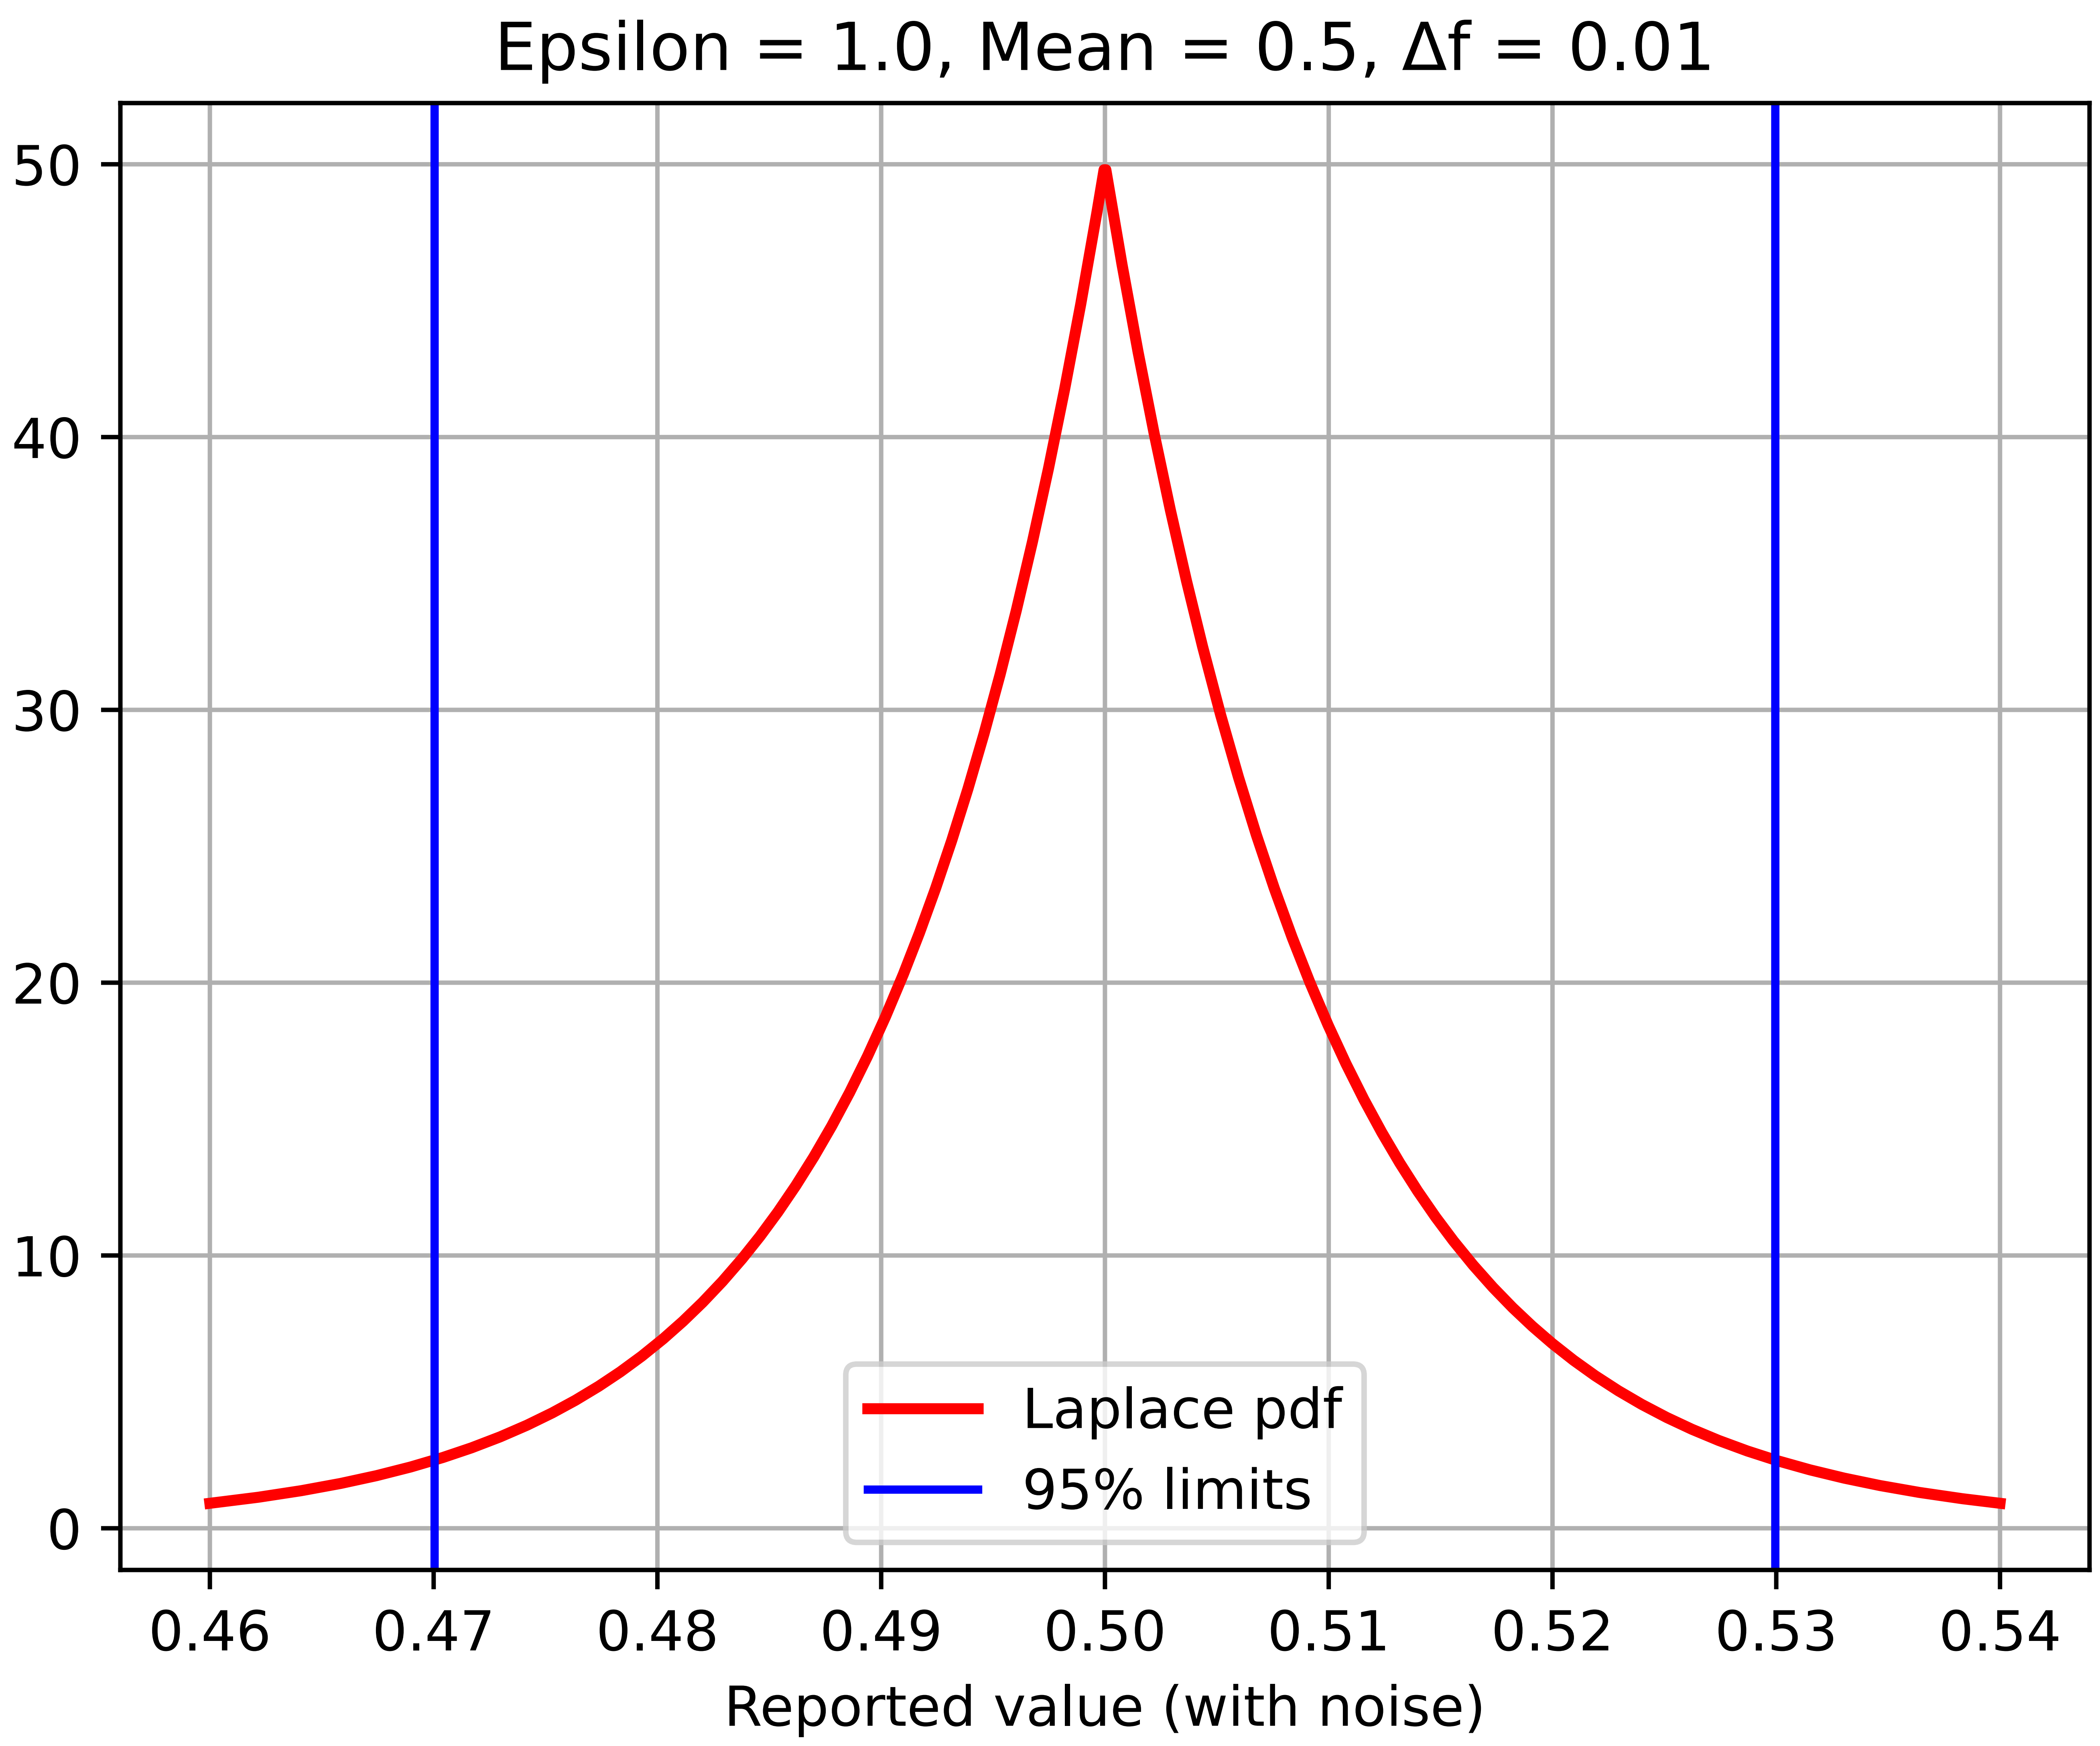

This is DP working as designed.

What if every company is 0% patched?

Binomial Pathology!!!

Company	# systems	% patched
Alpha	100	0
Bobble	100	0
Cantana	100	0
Delmax	100	0
…		
Company 49	100	0
Company 50	100	0
…		
Company 100	100	0

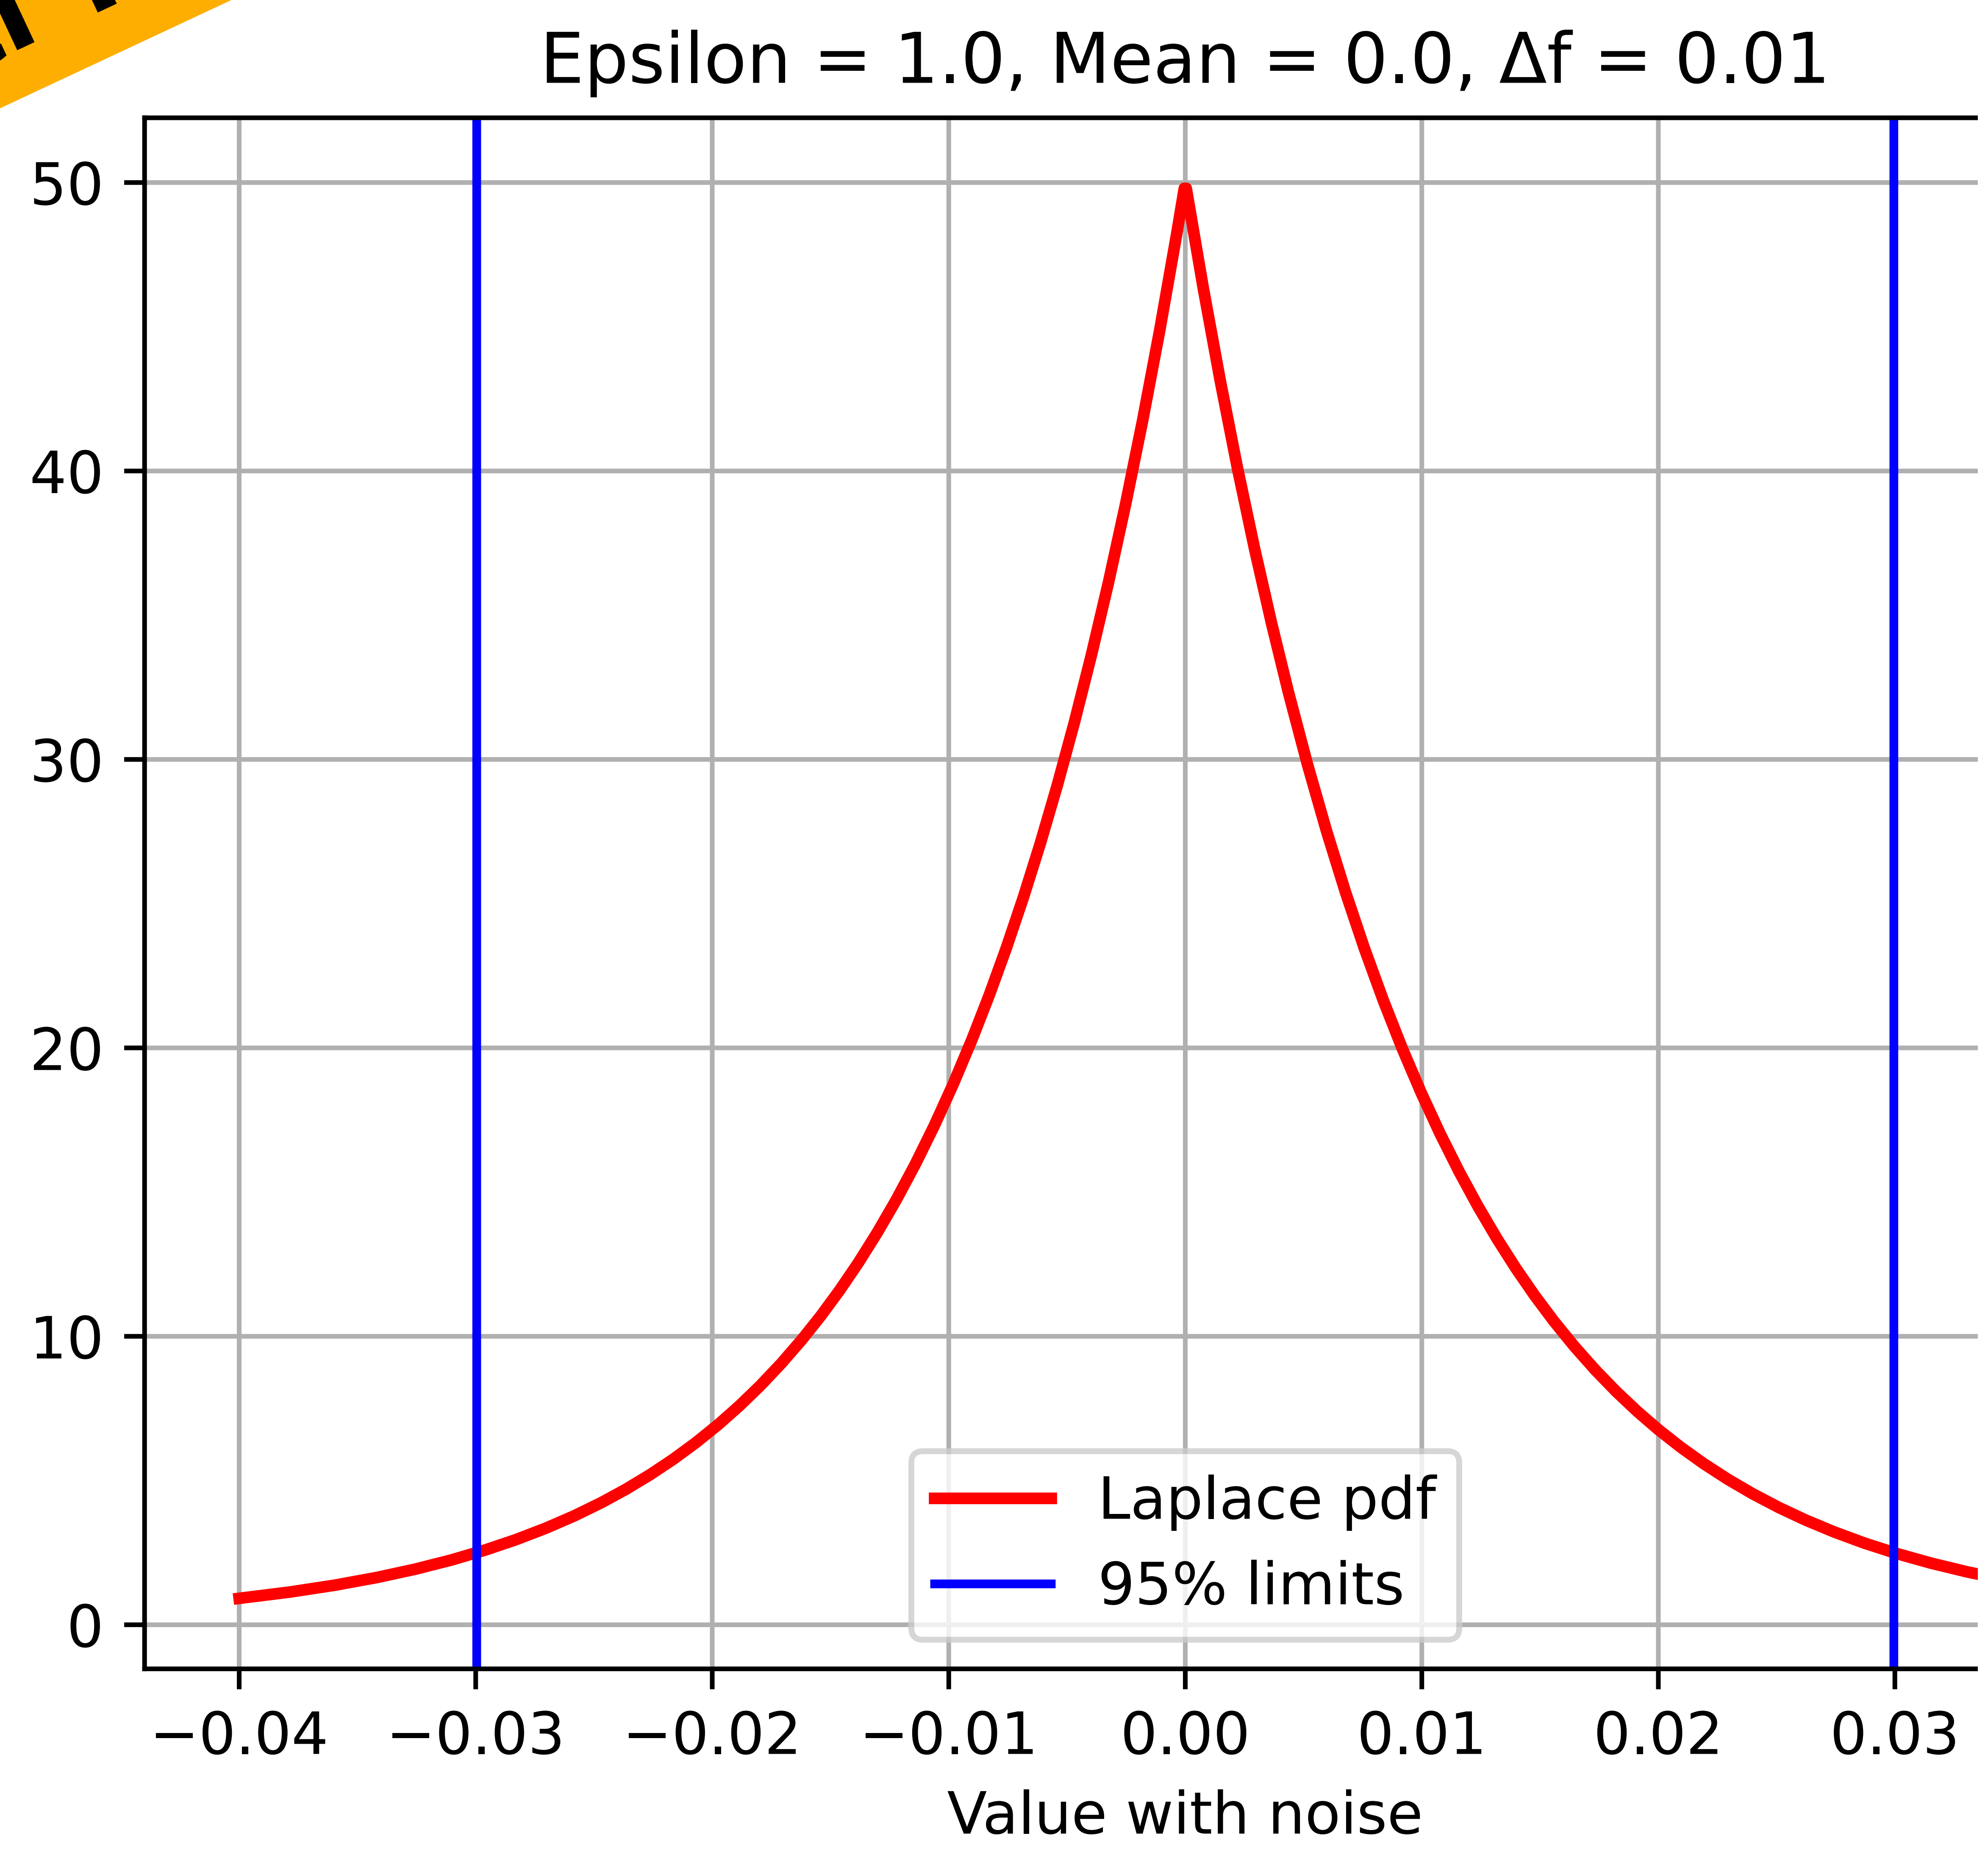

DP is not designed to protect this!

- Everybody looks equally bad!
- Even a company not included in the sample looks bad!
- How would you report the average is -2%?
- Notice these same problems happen if every company is 100% patched.

What if there are just 10 companies?

Company	# systems	% patched
Alpha	100	50
Bobble	100	50
Cantana	100	50
Delmax	100	50
Echo	100	50
Gulf	100	50
Hotel	100	50
Indigo	100	50
Julliet	100	50

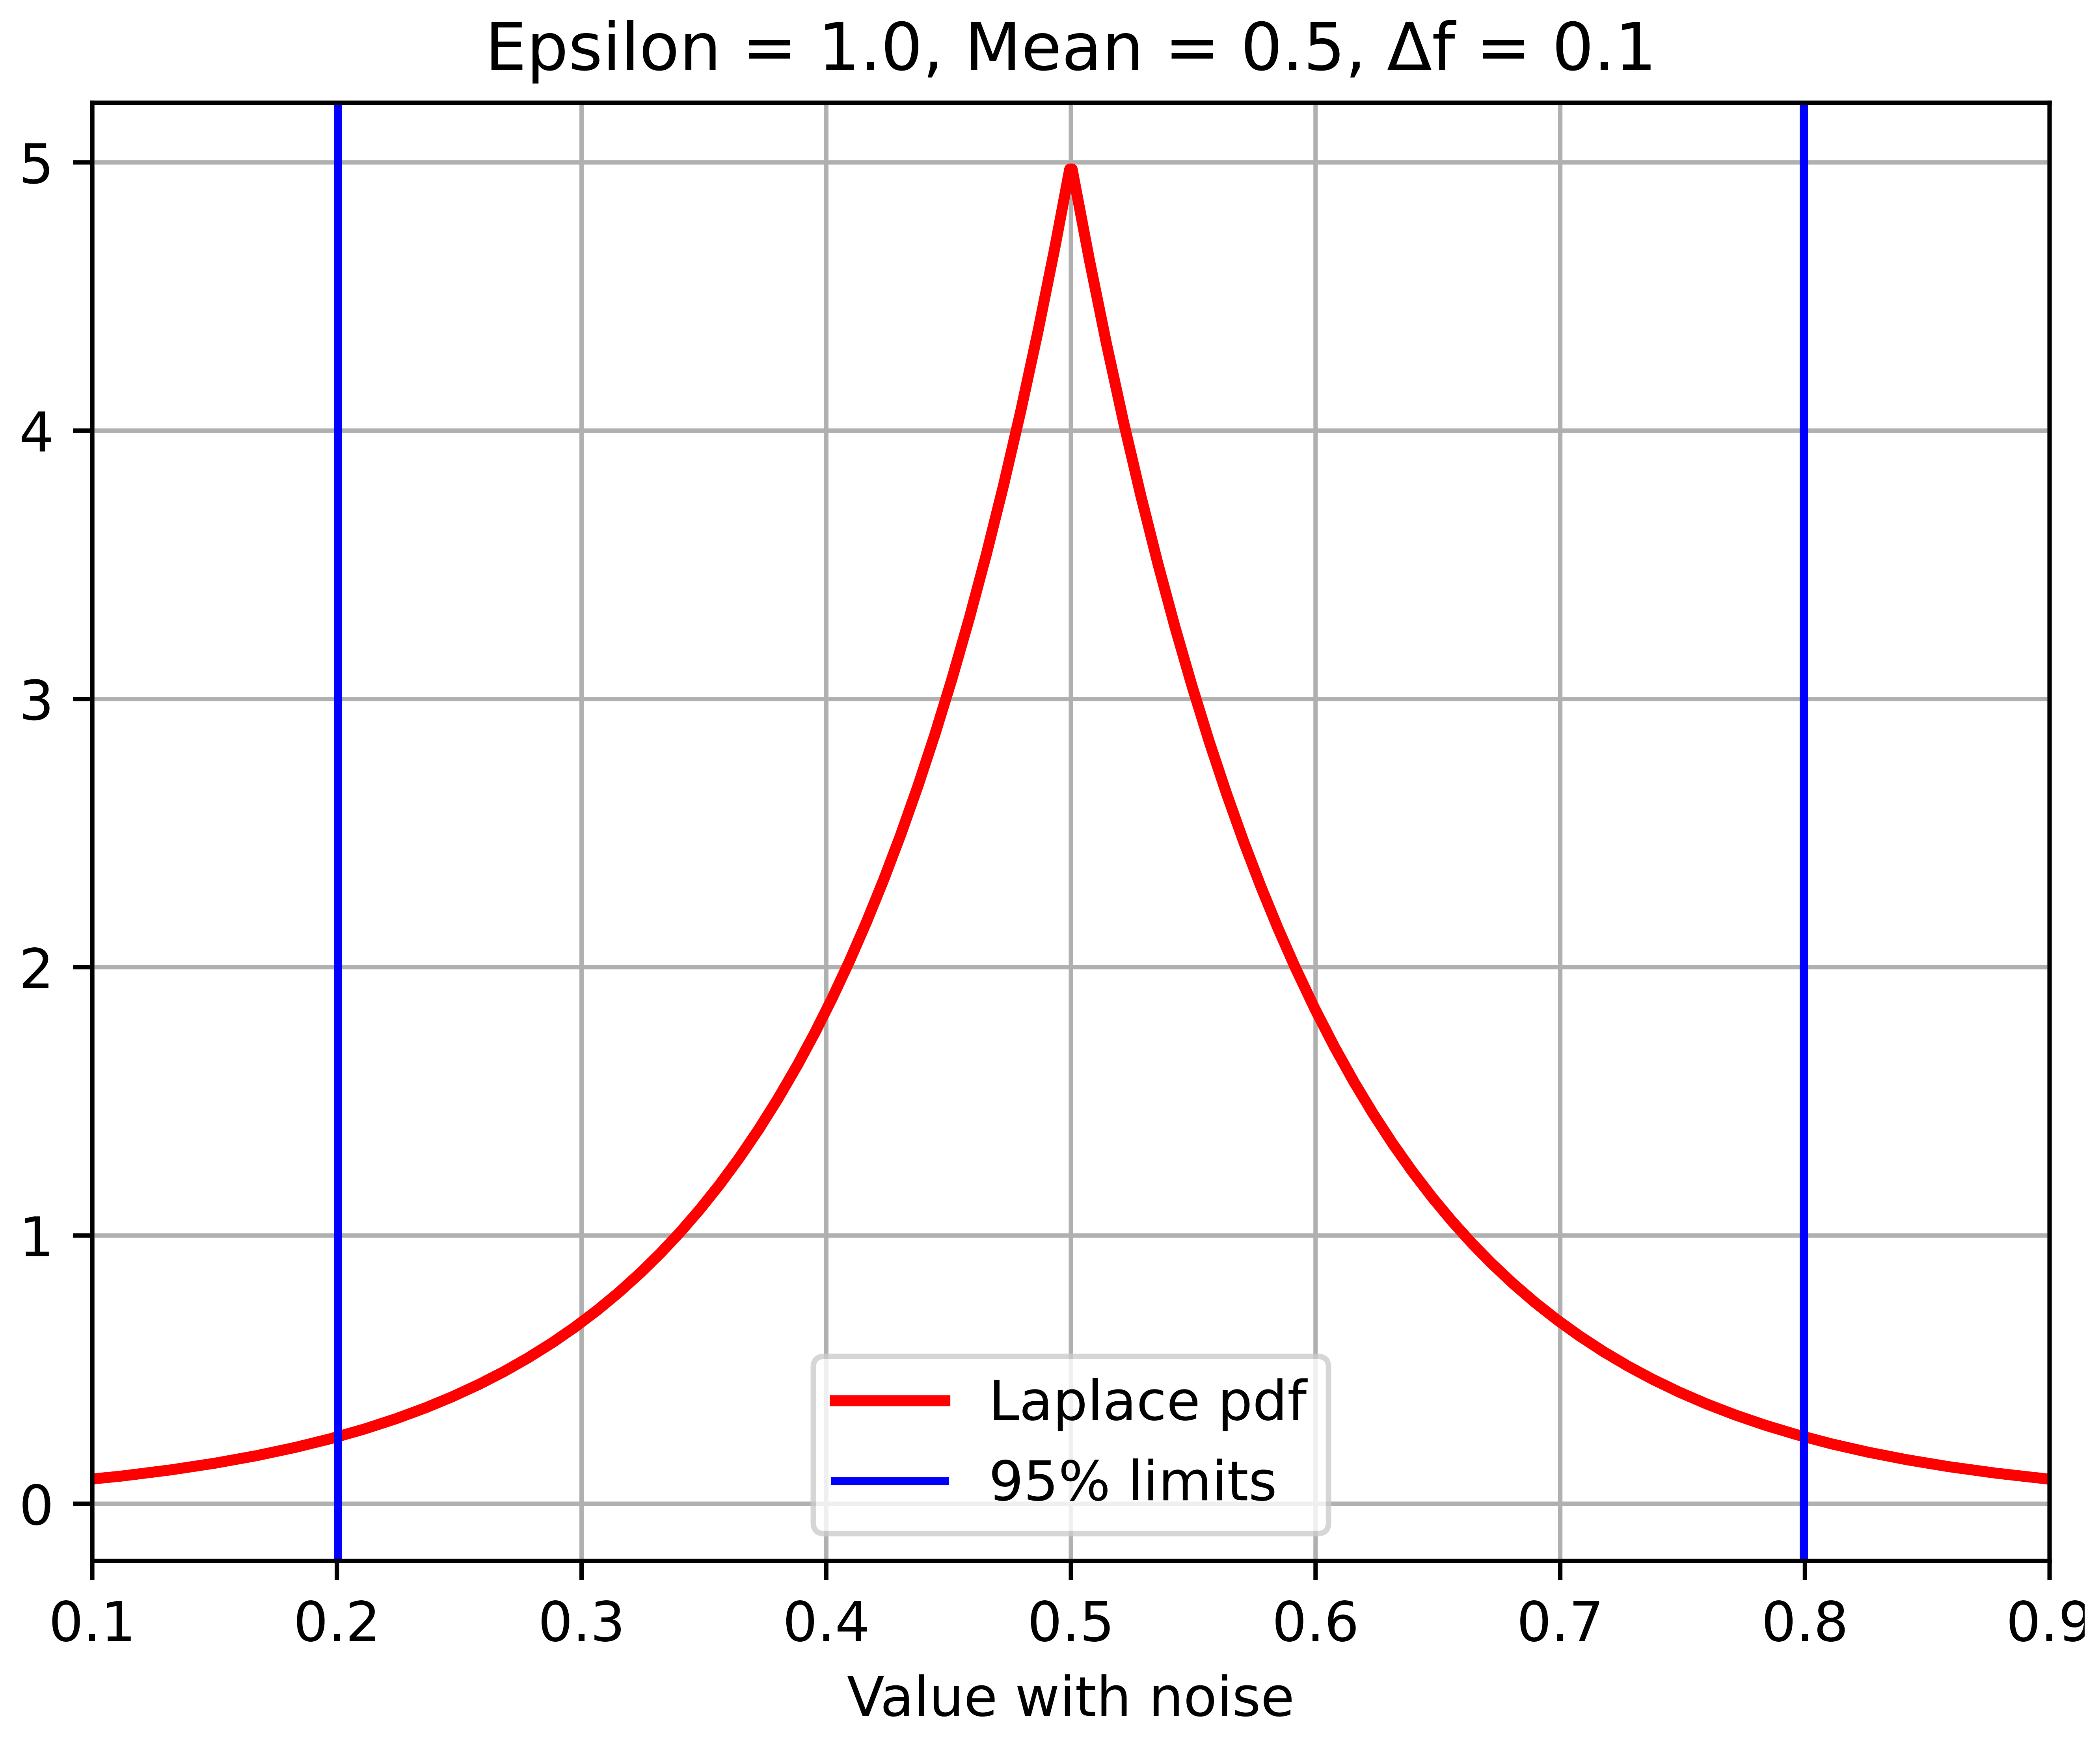

DP is designed to protect the worst case.
What if the attacker knows companies 1-9 are 50% patched?

Company	# systems	% patched
Alpha	100	50
Bobble	100	50
Cantana	100	50
Delmax	100	50
Echo	100	50
Gulf	100	50
Hotel	100	50
Indigo	100	50
Julliet	100	?

Binomial Pathology!!!

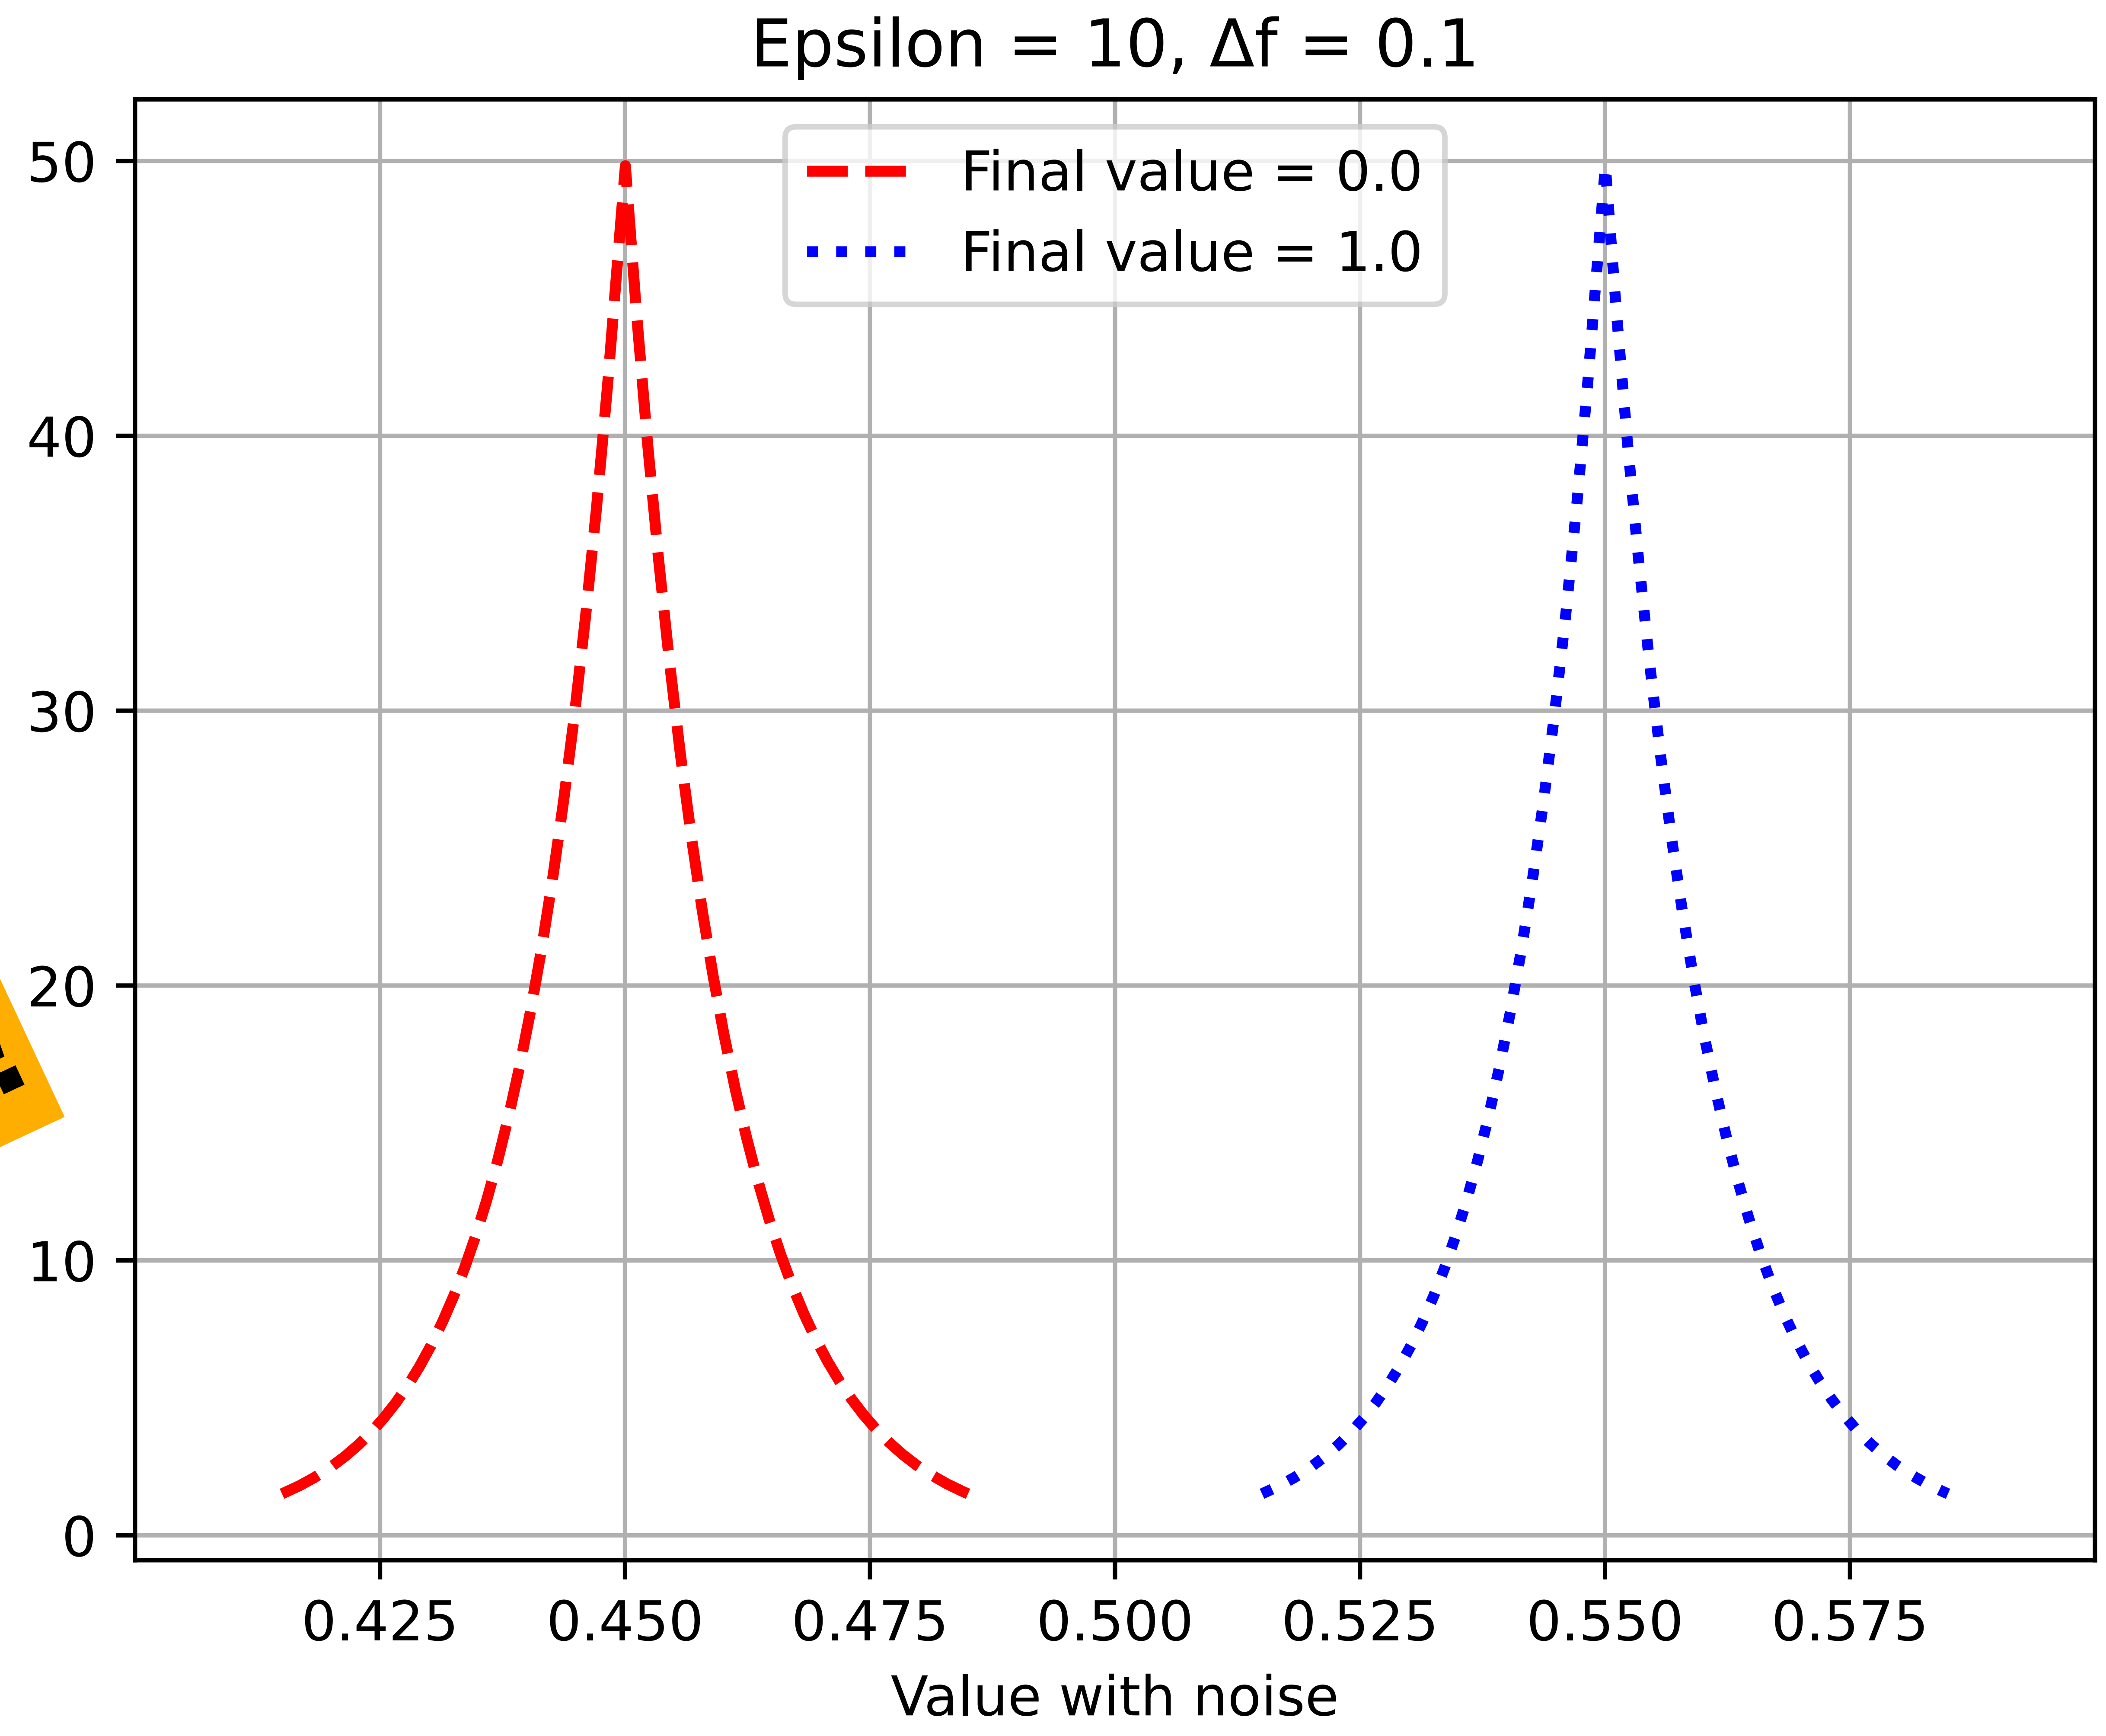

Now the attacker can get a good idea of company #10, at least with ε=1

Conclusions

Our focus is on harms, not the mathematical loss of privacy.

(in the paper)
We argue for a pragmatic (but thus risky) approach to adding noise.

Thank you!

David Clark — ddc@mit.edu
Simson Garfinkel — simsong@alum.mit.edu
KC Claffy — kc@sdsc.edu