Measuring the Loss of Privacy from Statistics

Michael Carl Tschantz
Carnegie Mellon University
Pittsburgh, PA, USA

Aditya V. Nori
Microsoft Research India
Bangalore, India

Meat

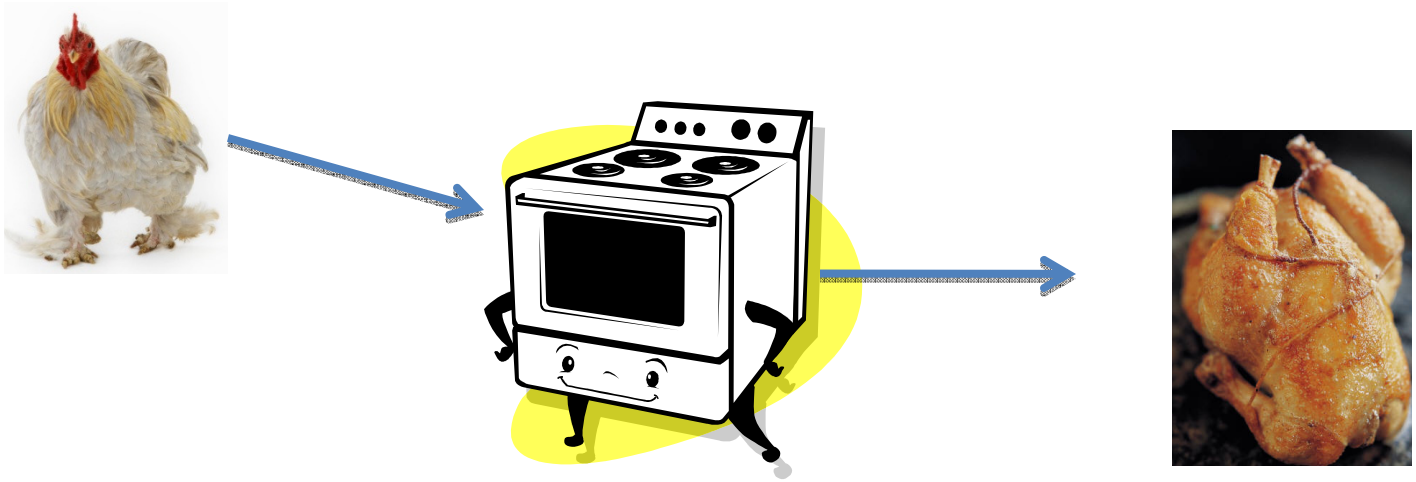

Meat

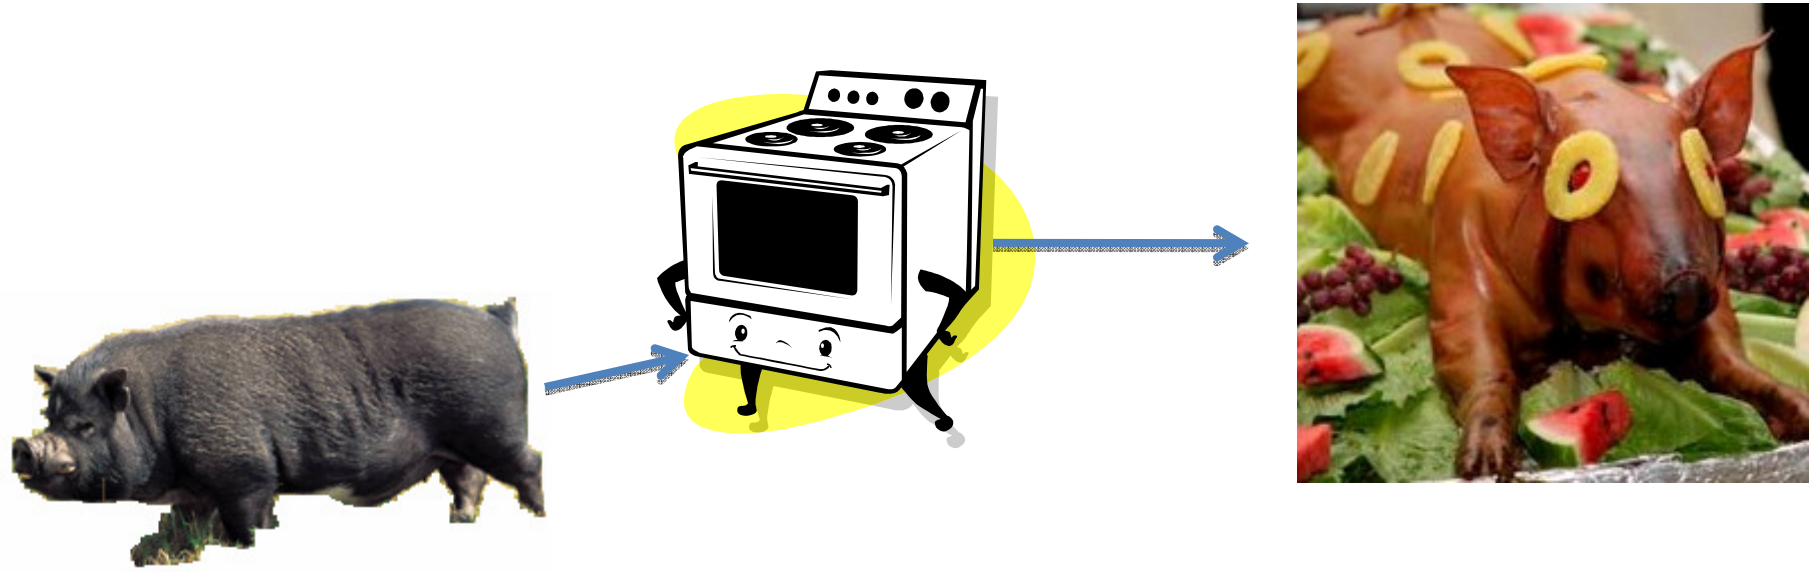

Meat

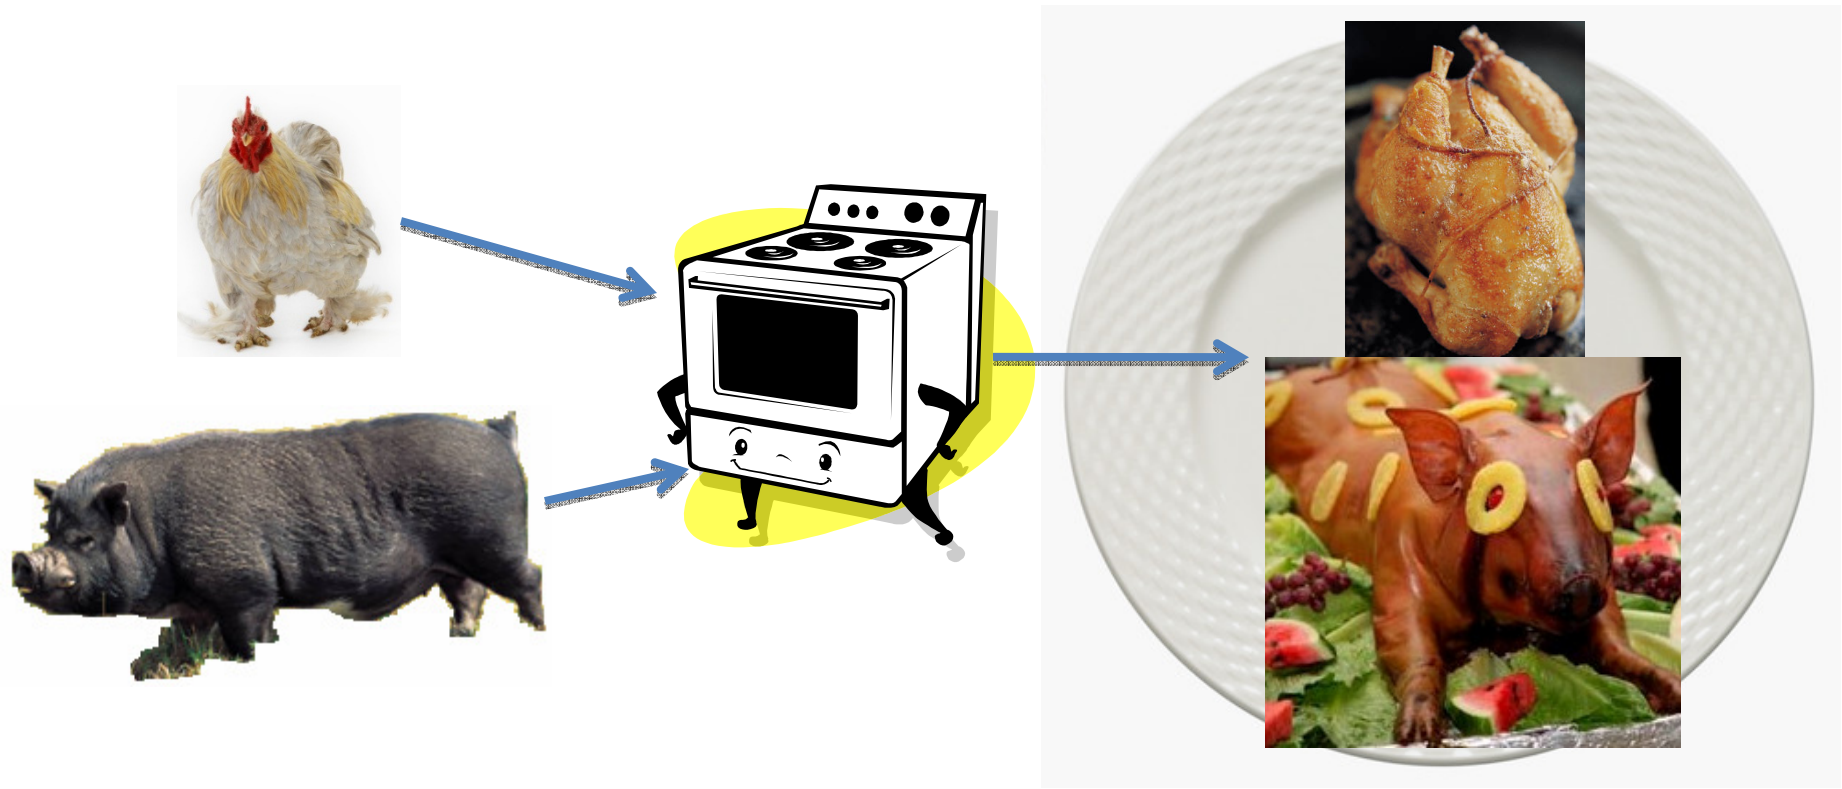

Meat

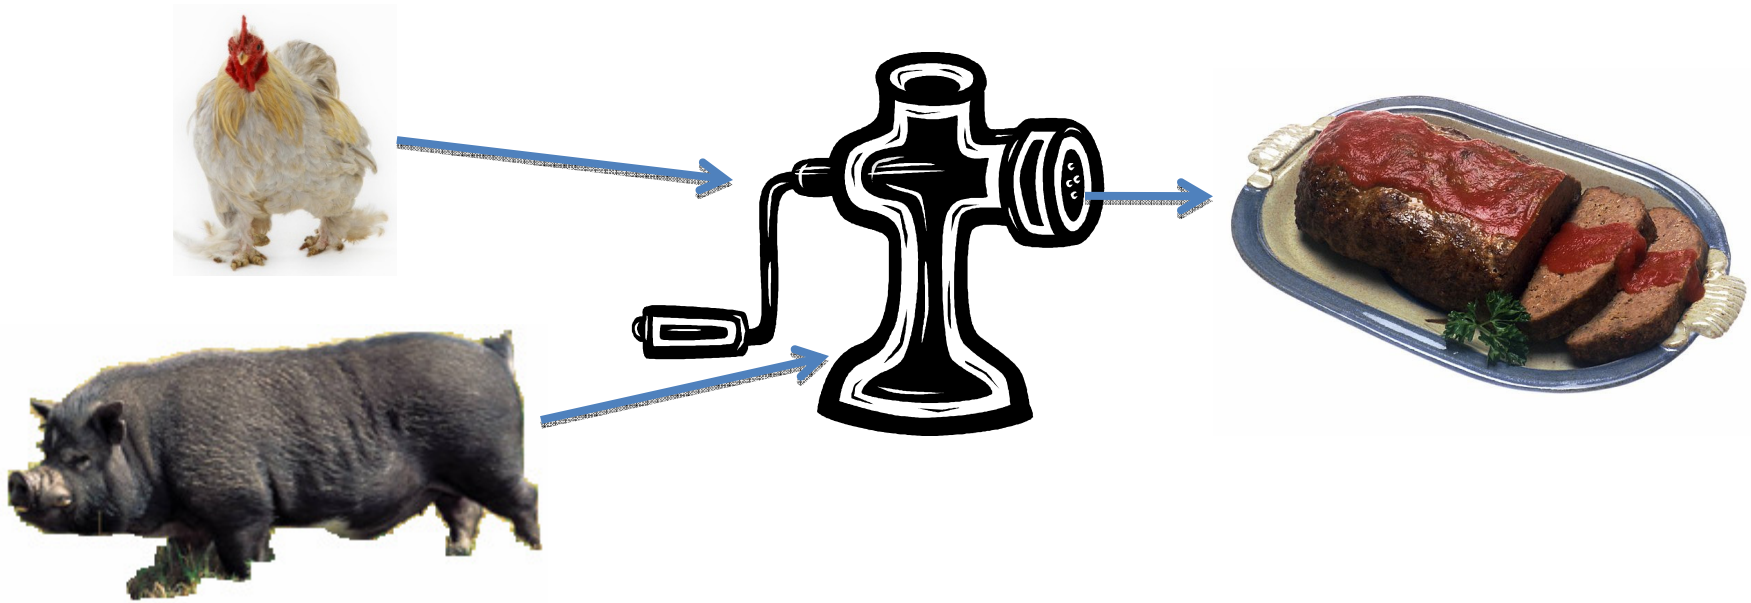

Meat

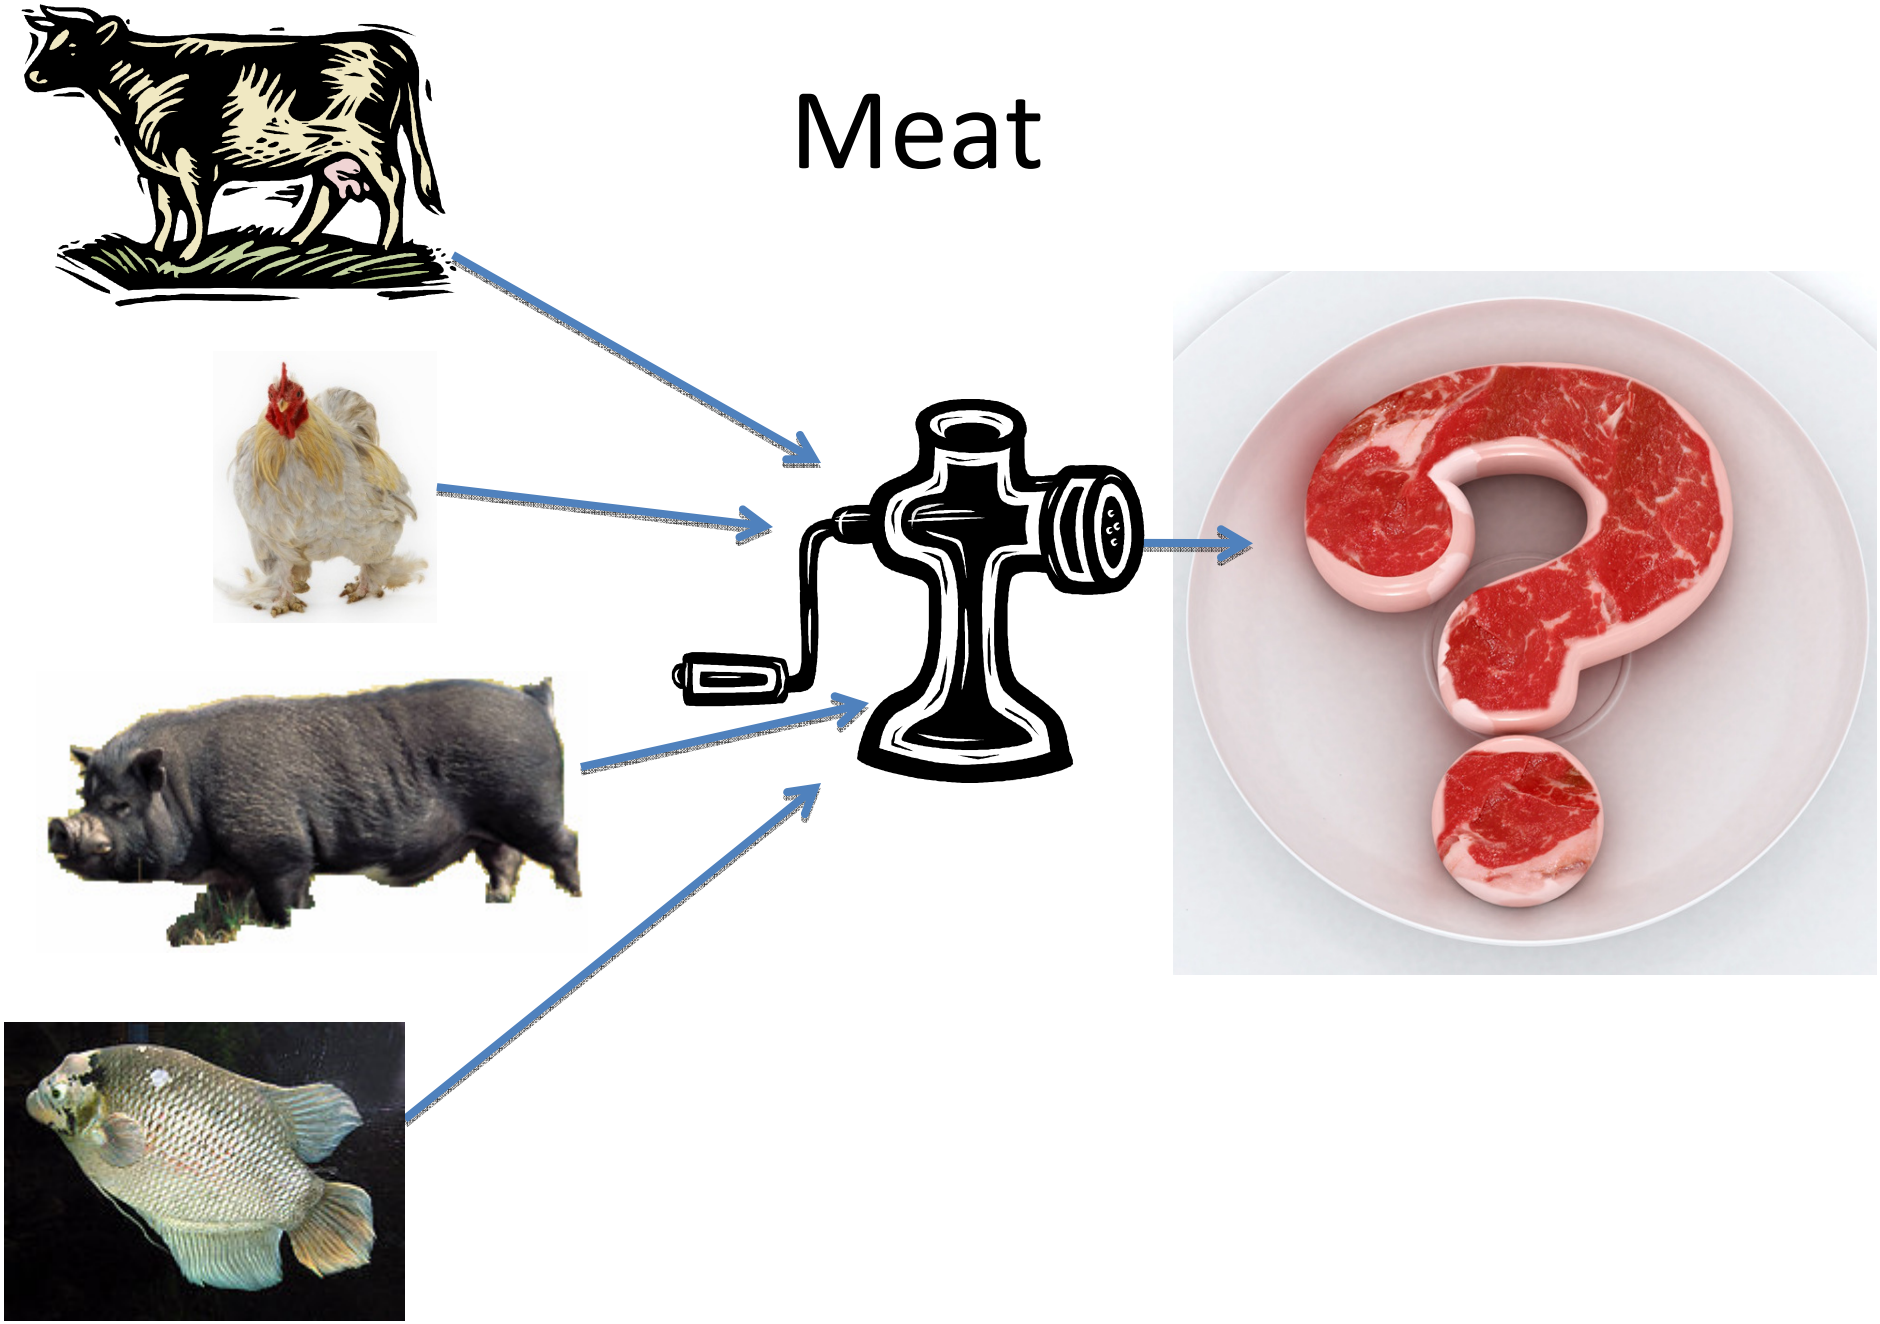

Information

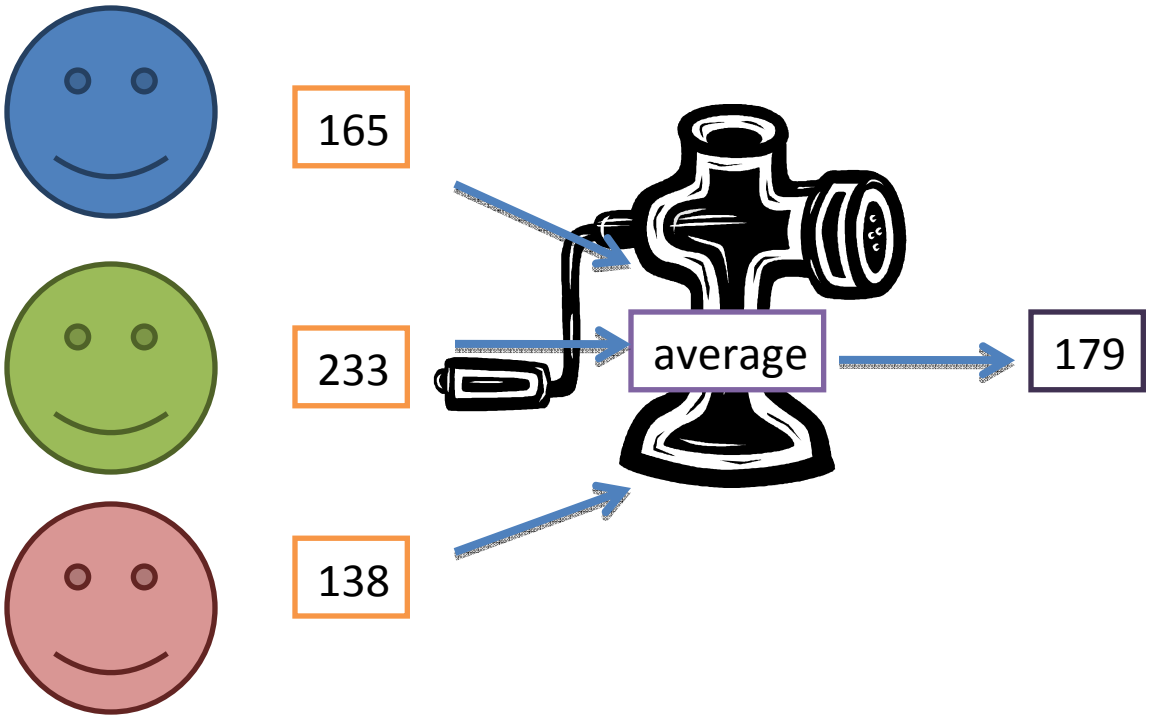

What does the average “179” tell us about Blue’s weight?

Information

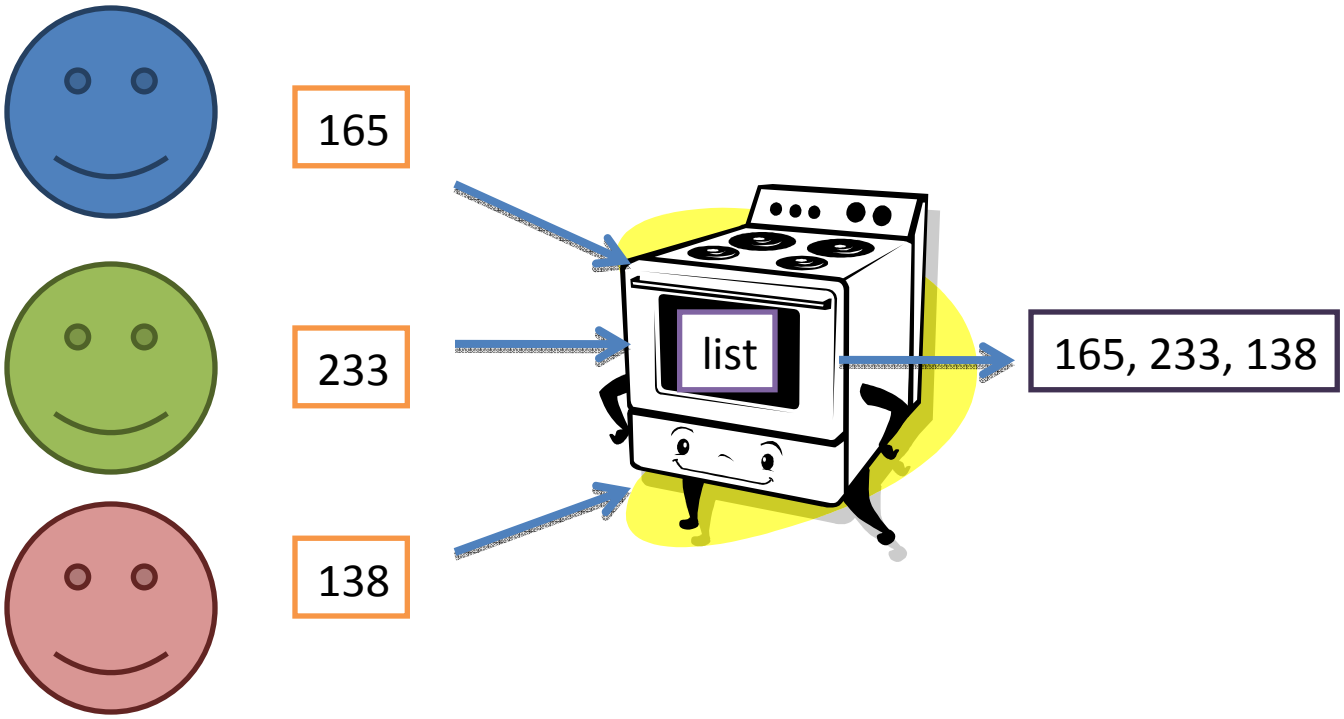

What does the list “165, 233, 138” tell us about Blue’s weight?

Adversary

- An adversary attempts to learn private info
  - What is Blue's weight?
- Has some prior beliefs
  - About 130
- Updates them based on statistic
  - An average of 179?! People weigh more than I thought

Goal

- Given program that computes statistics about a list of survey responses, characterize how much information the statistic provides about what an adversary is attempting to learn

Formal Model

- Program is a random variable **STAT** from survey responses to a statistic
- **STAT** takes on the actual value of **stat** based on the actual survey responses
- What the adversary would like to know is the value of a random variable **ADV**
- The adversary has prior beliefs **P** about the value that **ADV** takes on

Problem Statement

- Given
  - **STAT** program that computes statistics
  - **stat** the value the statistic takes on
  - **ADV** what the adversary wants to know
  - **P** an adversary's prior beliefs
- Compute the distribution for **ADV** under **P** given **STAT**=**stat** and compare to **ADV** under **P**

$$\mathbf{P}(\mathbf{ADV} \mid \mathbf{STAT}=\mathbf{stat}) \quad \text{vs.} \quad \mathbf{P}(\mathbf{ADV})$$

Measures of Change

- From the two distributions **P**(**ADV**) and **P**(**ADV** | **STAT**=**stat**) one can calculate:
  - Mutual information (change in entropy)
    H(**ADV**) – H(**ADV** | **STAT**=**stat**)
  - Change in Kullback–Leibler divergence
    Dist(**ADV**, **adv**) – Dist(**ADV**|**STAT**=**stat**, **adv**)
  - your favorite measure of distribution difference…

Adversary's Beliefs?

- Normally unknown
- From survey, we have an estimation of the actual probability distribution that produced the samples
- Use that in place of adversary’s beliefs to model an adversary that knows this distribution

Problem Statement

- Given
  - **STAT** program that computes statistics
  - **stat** the value the statistic takes on
  - **ADV** what the adversary wants to know
  - **P** an estimation of the under laying distribution
- Compute the distribution for **ADV** under **P** given **STAT**=**stat** and compare to **ADV** under **P**

  P(**ADV** | **STAT**=**stat**) vs. **P**(**ADV**)

Approach

- Monte Carlo Simulation
- Sample according to Prior-Beliefs **P**
- See how often **STAT** takes on the value **stat** for each value of **ADV**
- **P**(**ADV=adv** | **STAT**=**stat**)
  = **P**(**AVD=adv** & **STAT=stat**) / **P**(**STAT=stat**)
  ≈ #(**AVD=adv** & **STAT=stat**) / #(**STAT=stat**)

Performance

- The more samples, the more accurate
- Time linear in the number of samples and the amount of time **STAT** takes
- Memory linear in the range of **ADV** and memory usage of **STAT**

Convergence for Parity of X_1

- **H(Adv) – H(Adv|Stat=stat)** = 1 ≈ 0.999999797

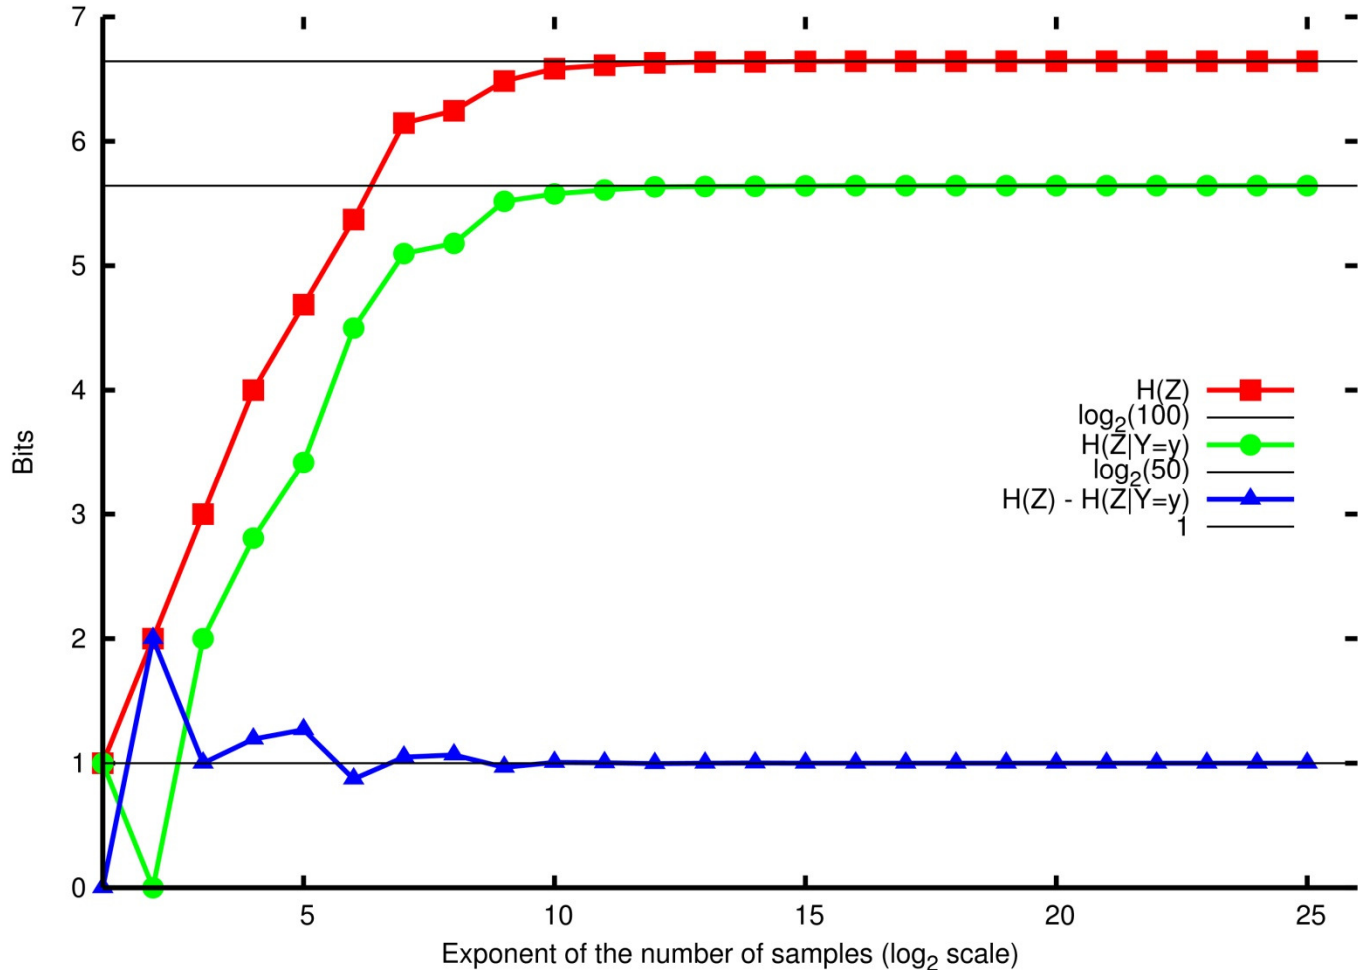

Mean, Median, Mode

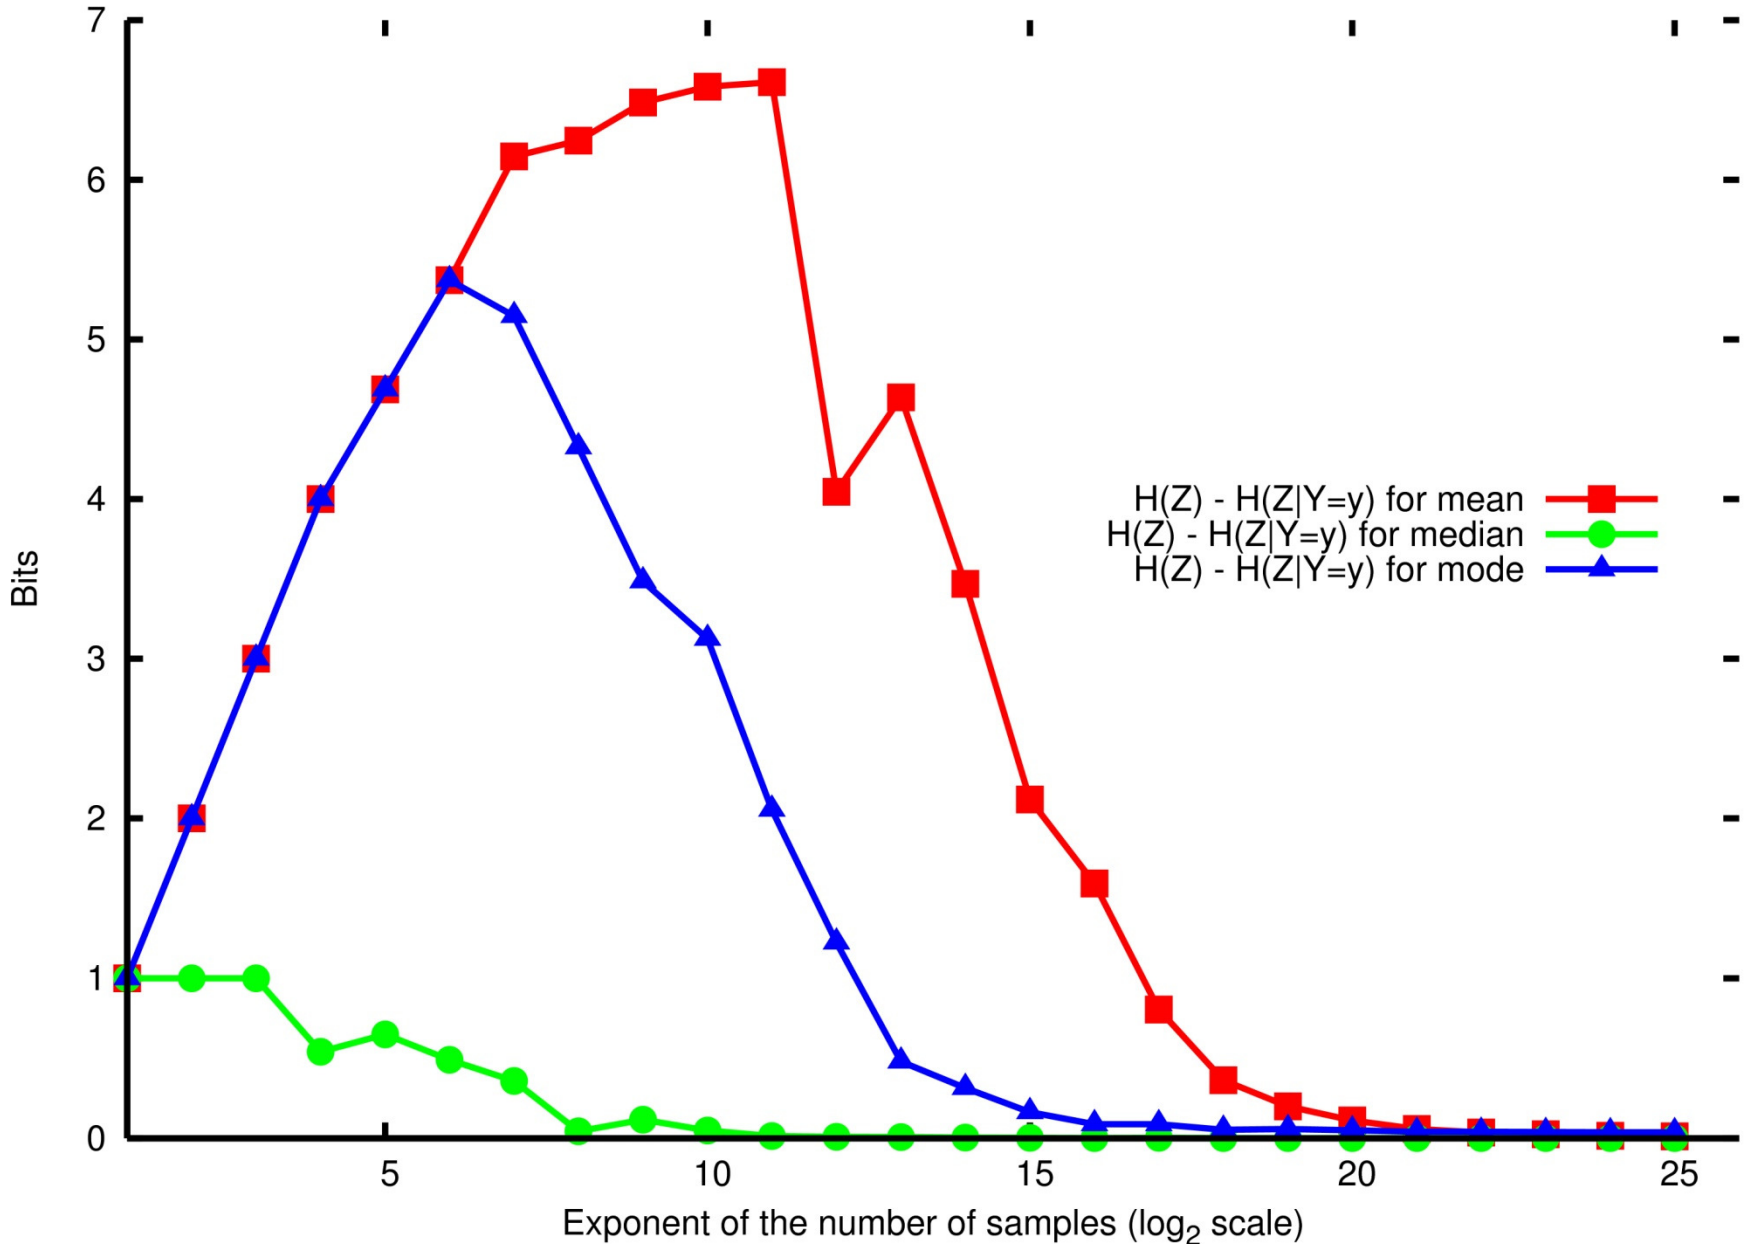

Comparison of Mutual Information

Statistic	**H(Adv) – H(Adv|Stat=stat)**	**Time (min)**
Parity	0.999999	10.65
Mean	0.012523	11.4
Median	0.002059	24.97
Mode	0.037691	40.73

Related Work:
Analyses for Mutual Information

- Mutual information is not always enough

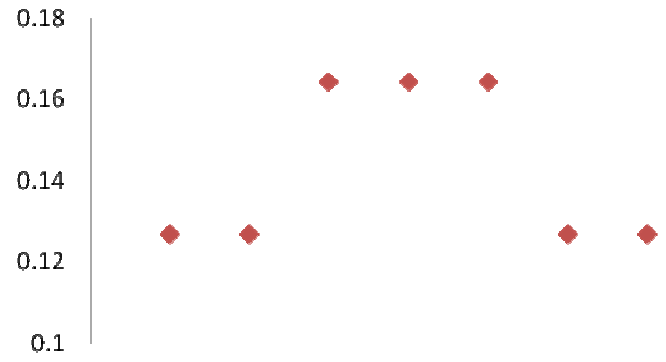

Related Work:
Analyses for Mutual Information

- Clark, Hunt, Malacaria – Static analysis
- McCamant and Ernst – Dynamic analysis
- Not exact enough (stove = meat grinder)

- Newsome and Song – Dynamic analysis
- Would be accurate enough with a theorem prover that finds all solutions to a logical formula

Related Work

Clarkson, Myers, Schneider

- Theory using adversary's beliefs
- No implementation
- Could be implemented using our work given adversary's beliefs

Related Work: Differential Privacy

- Dwork et al.
- Adds noise to protect privacy
- Does not distinguish between deterministic programs (stove = meat grinder)

Future Work

- Doesn’t work for really large sample spaces
- Doesn’t work if **STAT** is slow
- Modeling prior knowledge **P** is hard
- What to use for **ADV**

Questions?

- Implementation: http://www.cs.cmu.edu/~mtschant/mcqif/

Mean varying Survey Size

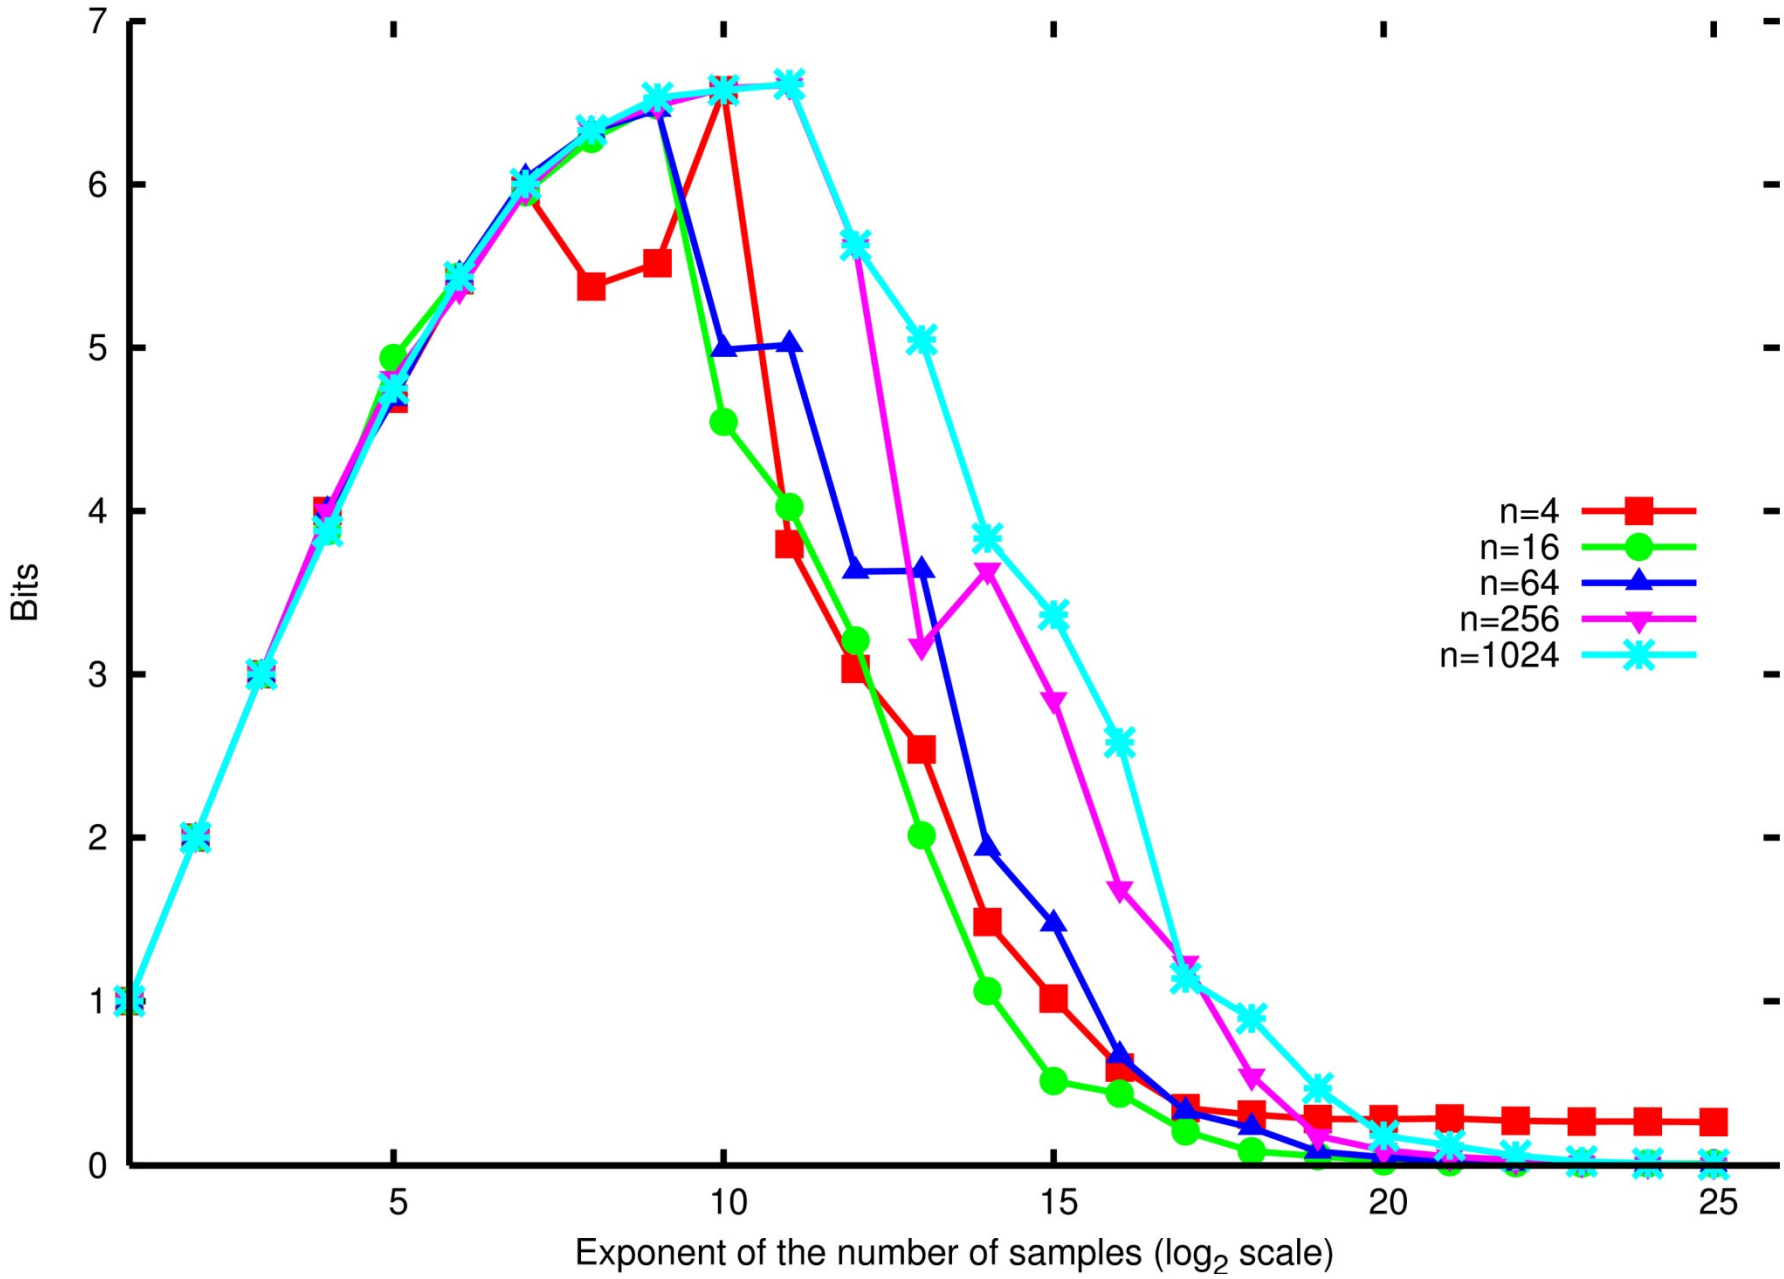

Related Work: Mutual Information

- D. Clark, S. Hunt, P. Malacaria
- Mutual info: H(**ADV**) – H(**ADV** | **STAT**=**stat**)
- Static system for measuring mutual info from single point to output
  - Not exact enough (stove = meat grinder)
  - Not always complete picture

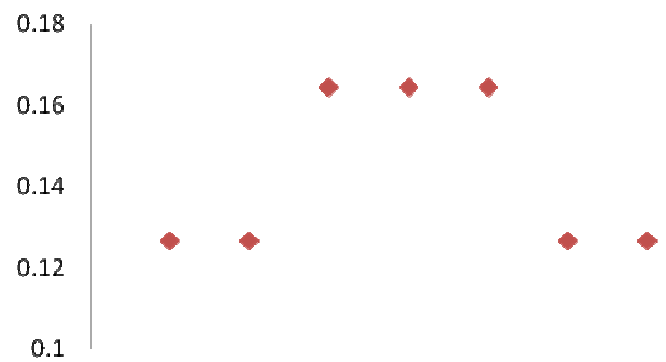

Change in Beliefs

- M. R. Clarkson, A. C. Myers, F. B. Schneider
- Dist(**ADV**, **adv**) – Dist(**ADV**|**STAT**=**stat**, **adv**)
  - Relative entropy
- Also not complete picture
- No implementation

Channel Capacity

- Newsome and Song
- Converts single execution trace of program to a logical formula
- Use theorem prover to find all solutions
  - Can only provide lower bound in practice
- Bounds information flow for that trace for any input distribution
  - We use a fixed input distribution