Efficient Data Collection Requires Incentives

MADLab 2022 Summer Workshop

Yiding Chen, Young Wu, Jerry Zhu*
Department of Computer Sciences, University of Wisconsin-Madison. June 14, 2022

A Fishy Tale

Single Agent Data Collection: Payoff Maximization

- Collect n data points $x_1, \ldots, x_n \sim N(\mu, 1)$
- Maximum Likelihood Estimate $\hat{\mu} = \frac{1}{n} \sum_{j=1}^{n} x_j$
- Variance $V(\hat{\mu}) = \frac{1}{n}$
- Key assumption: Agents collect data for economic reasons

Single Agent Data Collection: Payoff Maximization

- Benefit to the agent: $b(n)$ concave in n (diminishing return)
  - Example: $b(n) = \sqrt{1/V(\hat{\mu})} = \sqrt{n}$
- Cost to the agent: $c(n)$
  - Example: $c(n) = \alpha n$, α is unit data collection cost
- Payoff to agent $u(n) = b(n) - c(n)$

Single Agent Data Collection: Payoff Maximization

- How many fish would Rob want to measure?

$$\max_{n} u(n) = \sqrt{n} - \alpha n$$

$$\frac{1}{2\sqrt{n}} - \alpha = 0$$

$$n = \frac{1}{4\alpha^2}$$

Single Agent Data Collection: Payoff Maximization

$$n = \frac{1}{4\alpha^2}$$

"Continuous fish"

$$u(n) = \frac{1}{4\alpha}$$

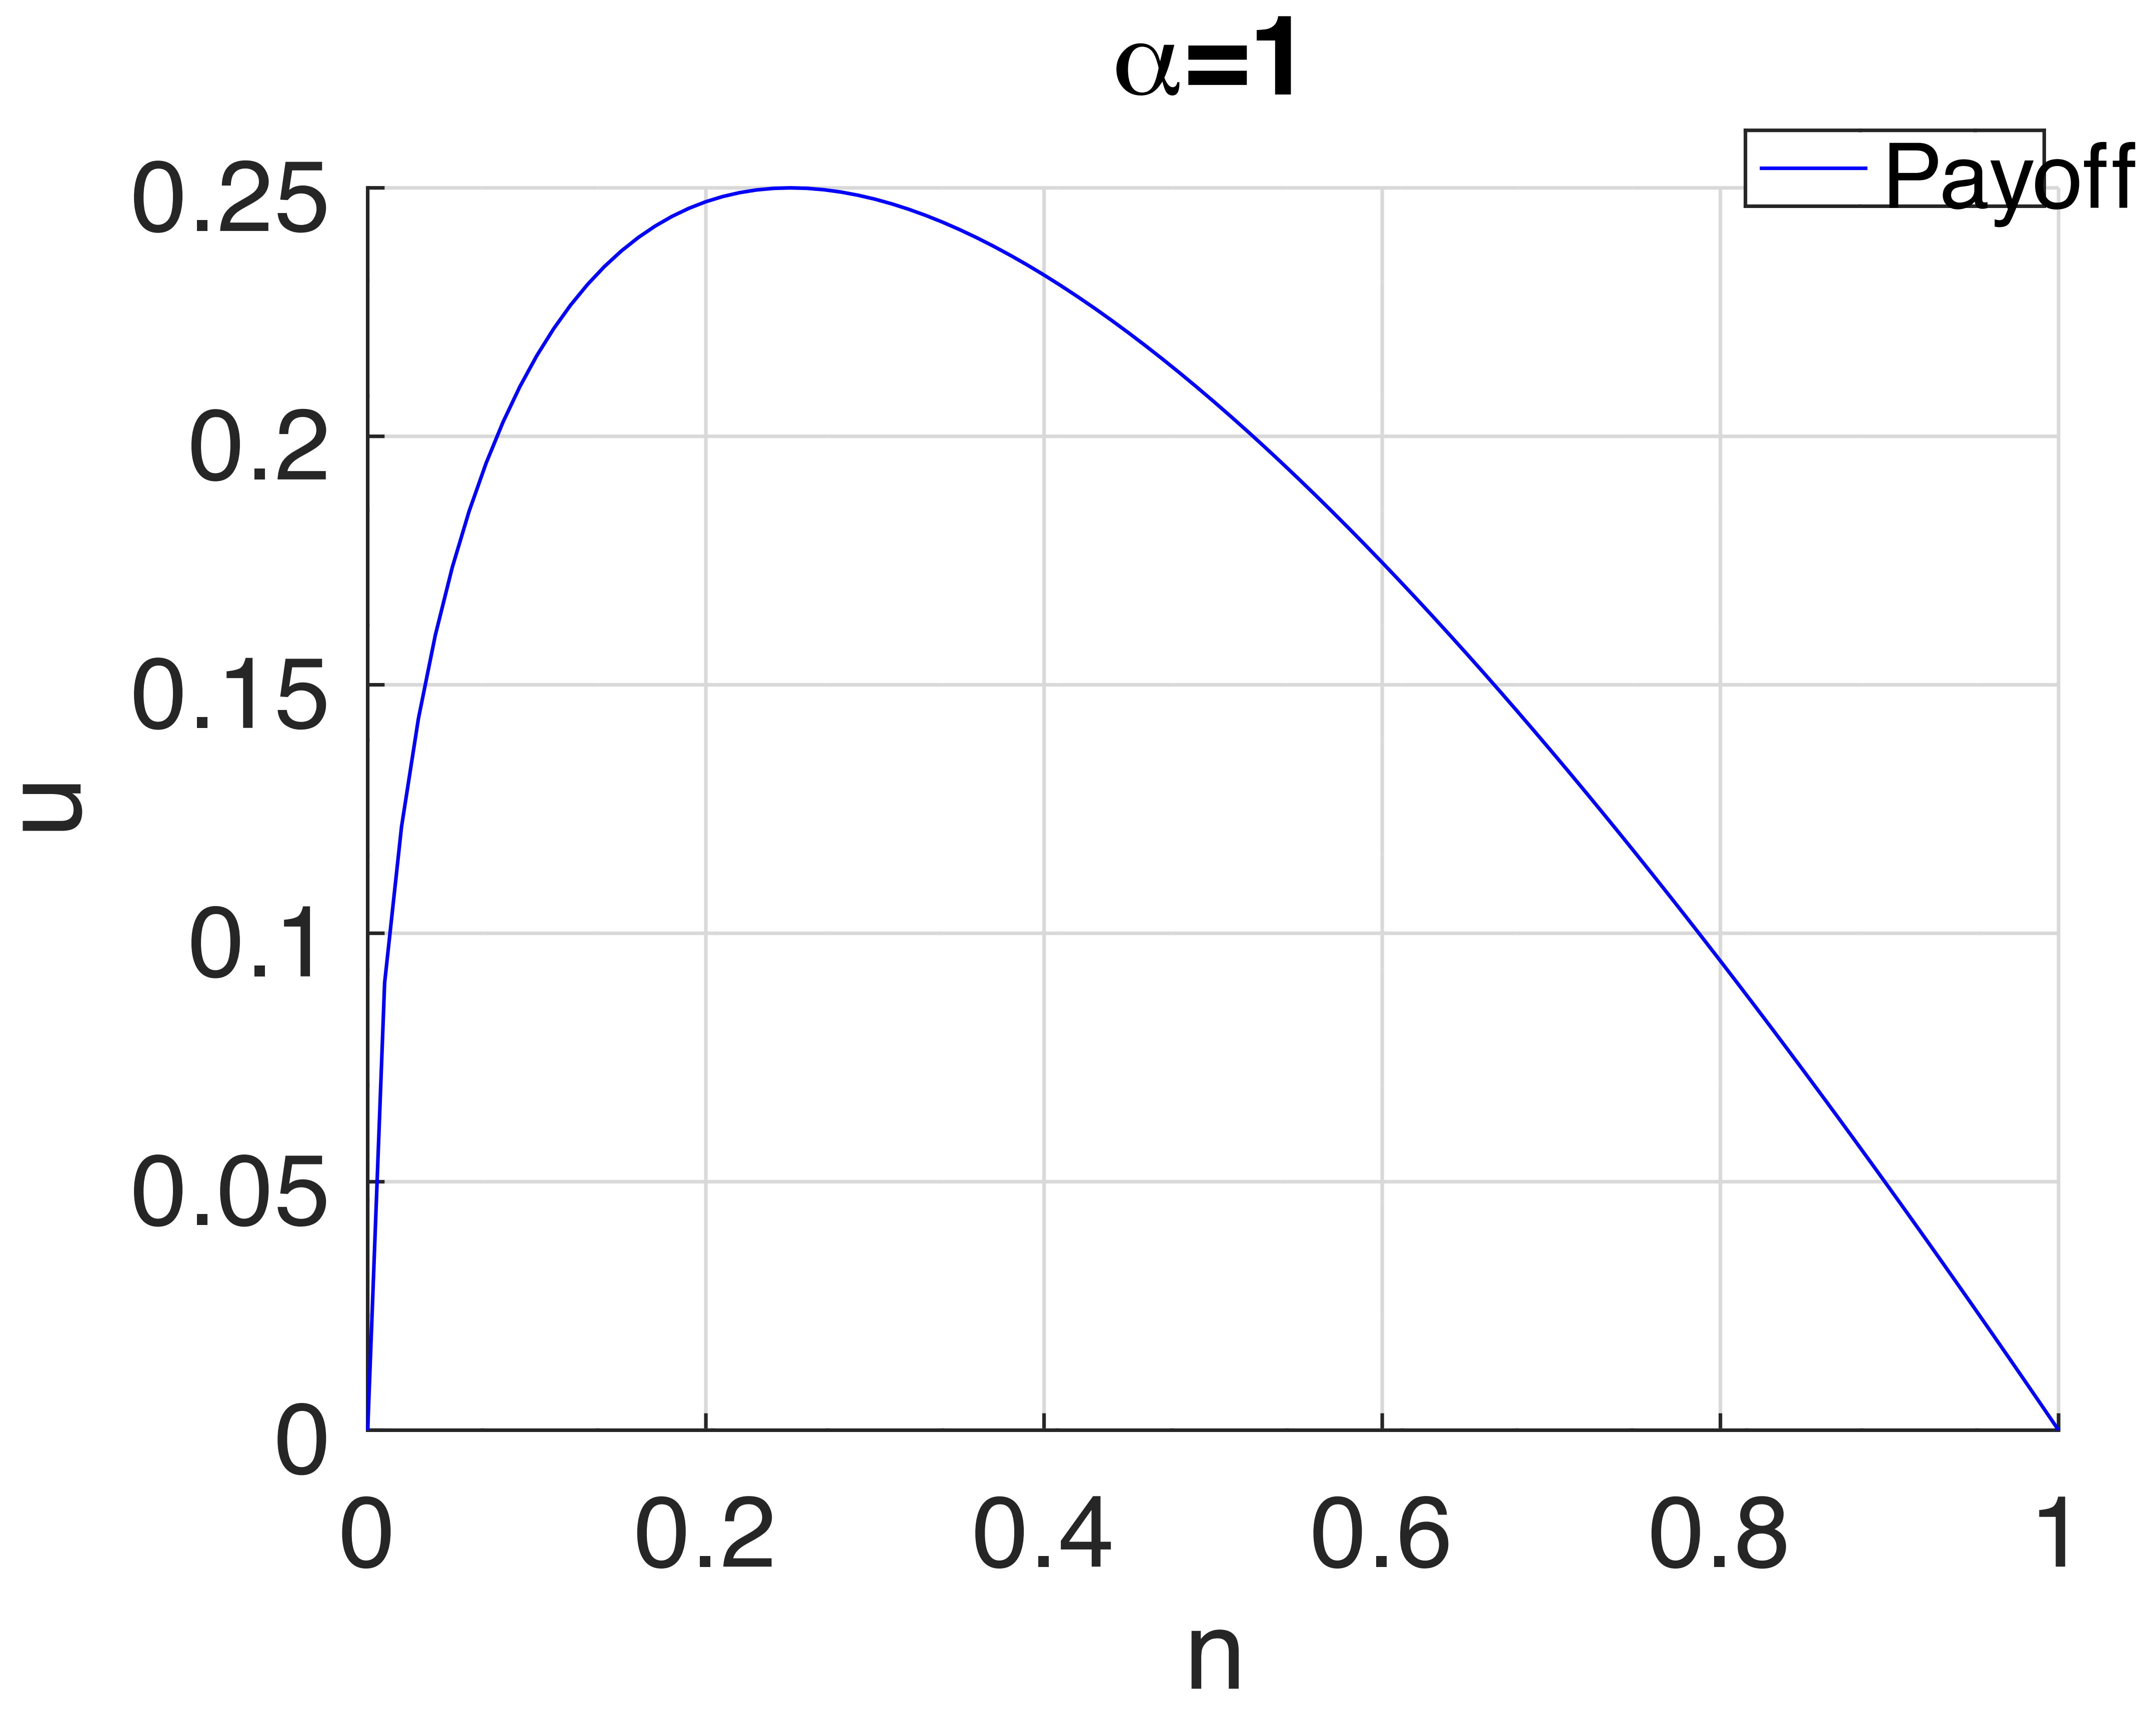

Let's All Measure Fish

- m autonomous agents, each collecting $n_1, \ldots, n_m$ iid points from $N(\mu, 1)$
- We will share our data
- $V(\hat{\mu}) = \dfrac{1}{\sum_{i=1}^{m} n_i}$
- Payoff to agent i:

$$u_i(n_1, \ldots, n_i, \ldots, n_m) = \sqrt{\sum_{j=1}^{m} n_j} - \alpha n_i$$

Rationality

- Agent i action space $n_i \in \mathbb{R}_{\geq 0}$
- The Best Response to other agents who play $n_{-i} = (n_1, \ldots, n_{i-1}, n_{i+1}, \ldots, n_m)$:

$$n_i^{BR} \in \arg\max_{n_i} u_i(n_i, n_{-i})$$

- "If others play n_{-i}, I do not want to deviate from n_i^{BR}."
- But what *will* others play?

Nash Equilibrium

- $(n_1^*, \ldots, n_m^*)$ is a Nash equilibrium if the components are BR to each other:

$$n_i^* \in \arg\max_{n_i} u_i(n_i, n_{-i}^*), \forall i \in [m]$$

- Our problem is symmetric
- *Assume* a symmetric NE $(n^*, \ldots, n^*)$

Multi-Agent Rational Data Collection

$$n_i^* \in \arg\max_{n_i} u_i(n_i, n_{-i}^*), \forall i \in [m]$$

$$\left.\frac{\partial u_i(n_i, n_{-i} = n^*)}{\partial n_i}\right|_{n_i = n^*} = 0$$

$$n^* = \frac{1}{4\alpha^2 m}$$

Multi-Agent Rational Data Collection

$$n^* = \frac{1}{4\alpha^2 m} \qquad u_i(n^*, \ldots, n^*) = \frac{1}{2\alpha} - \frac{1}{4\alpha m}, \forall i \in [m]$$

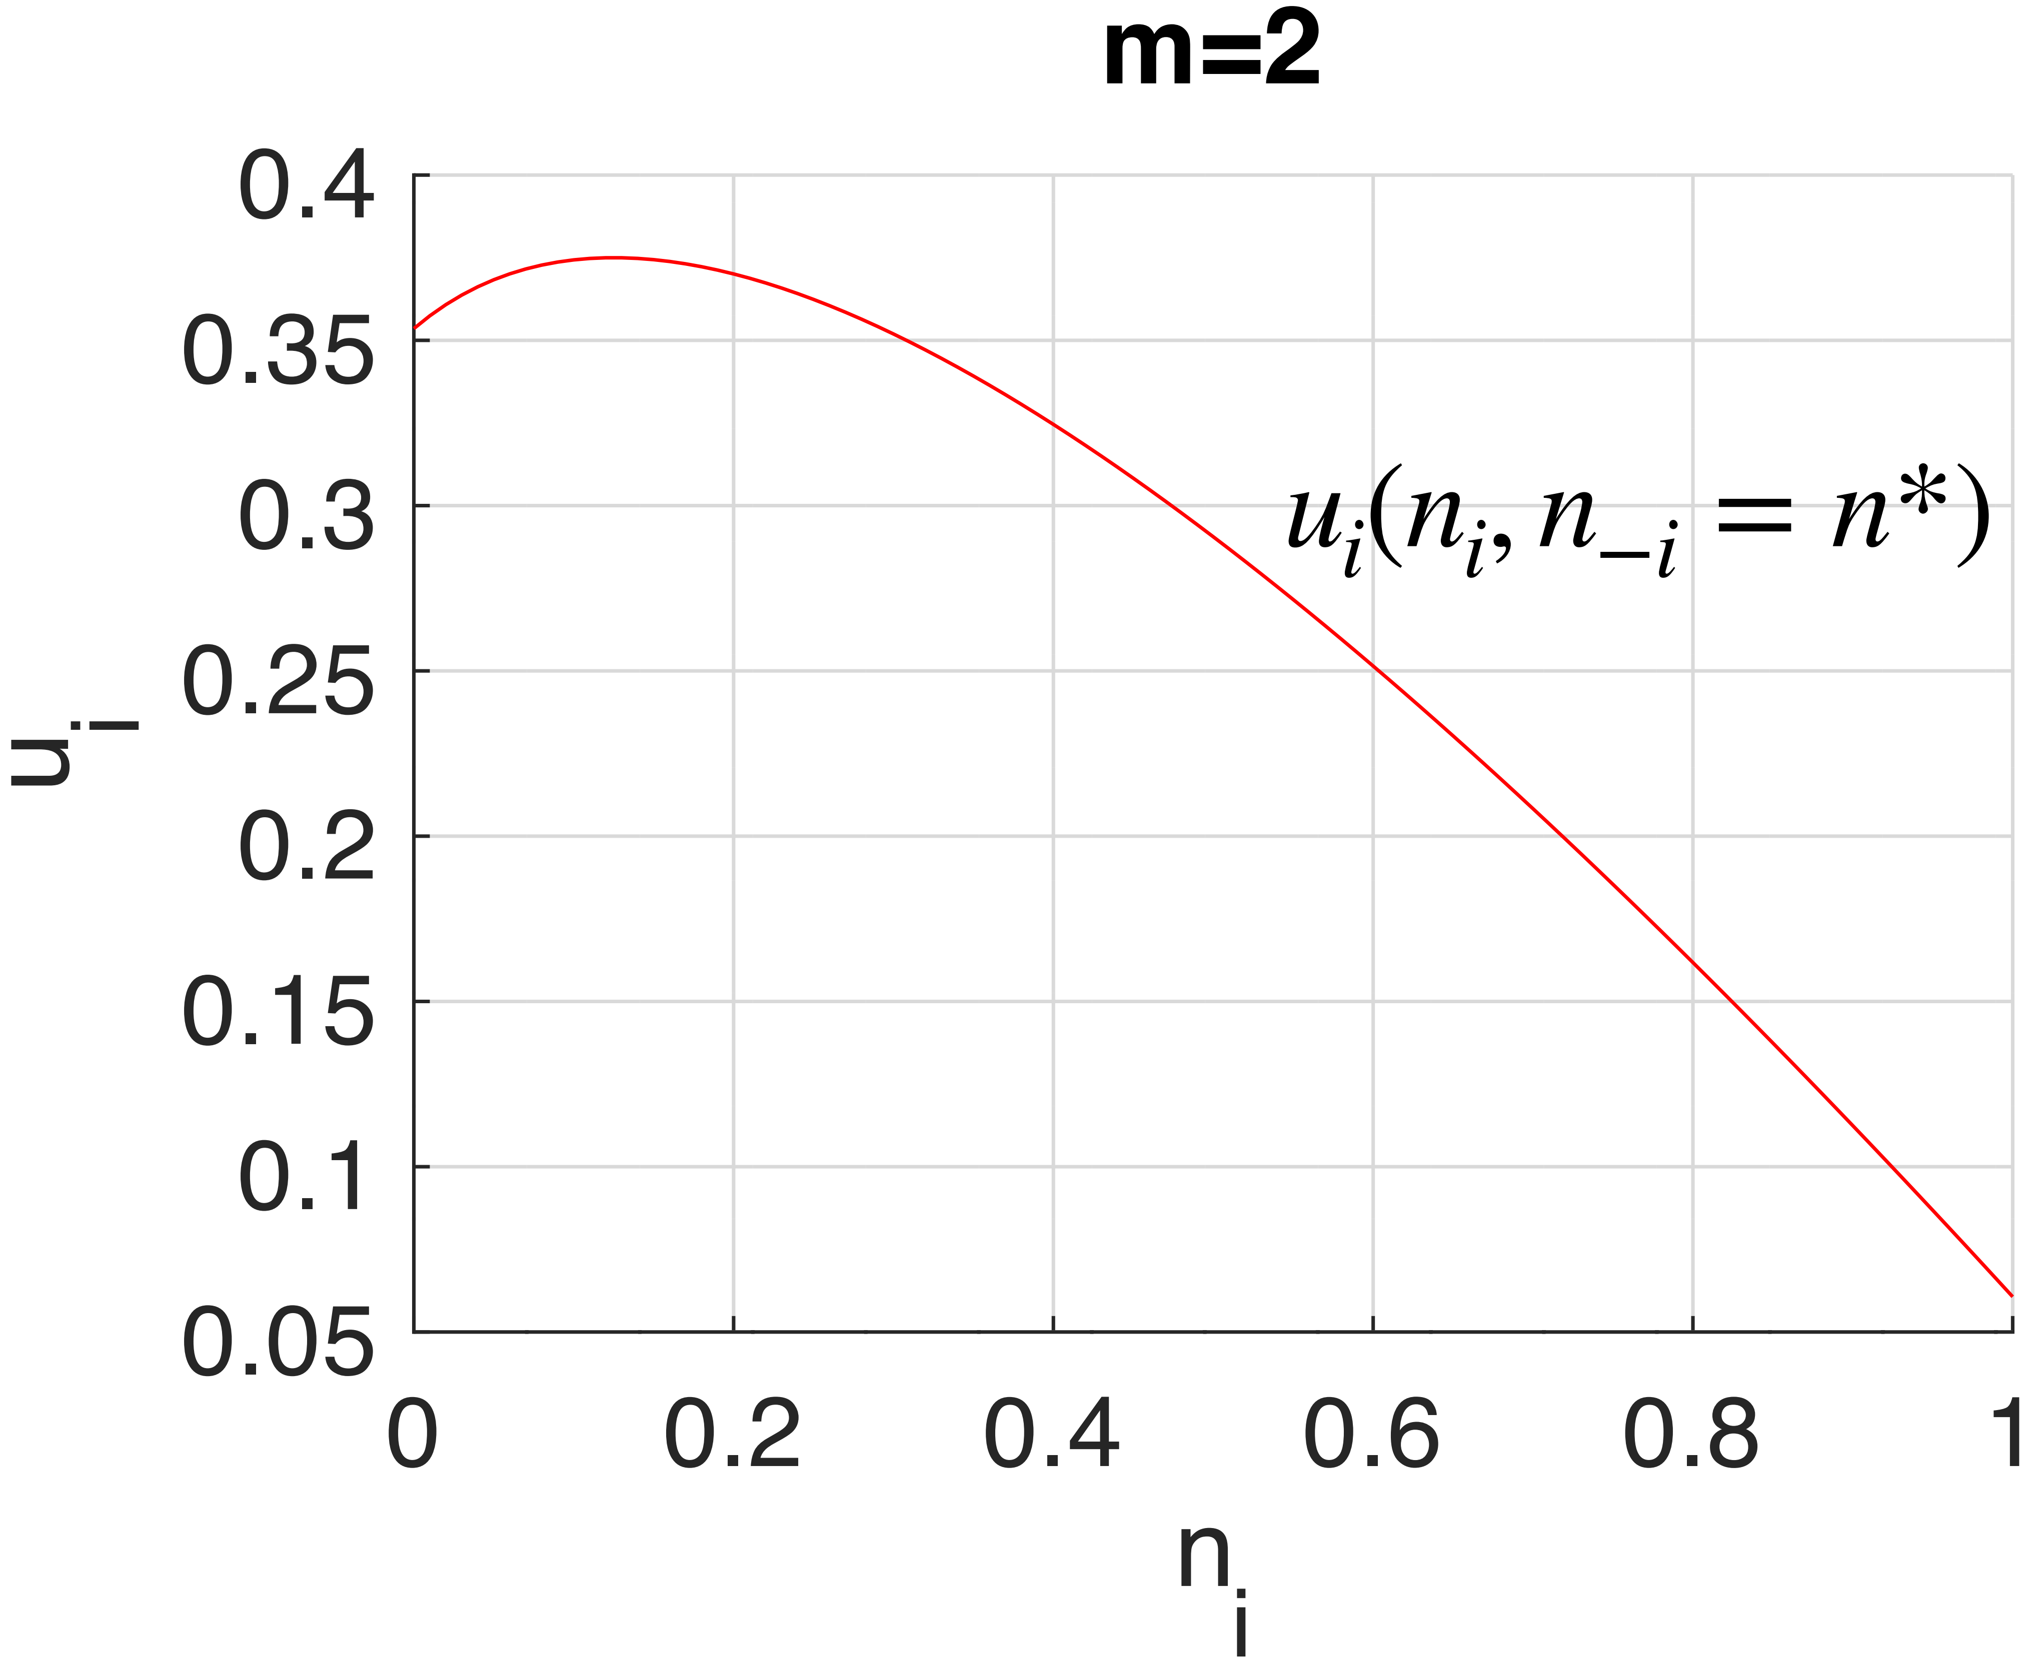

Compared to Rob alone

$$n = \frac{1}{4\alpha^2}$$

$$u(n) = \frac{1}{4\alpha}$$

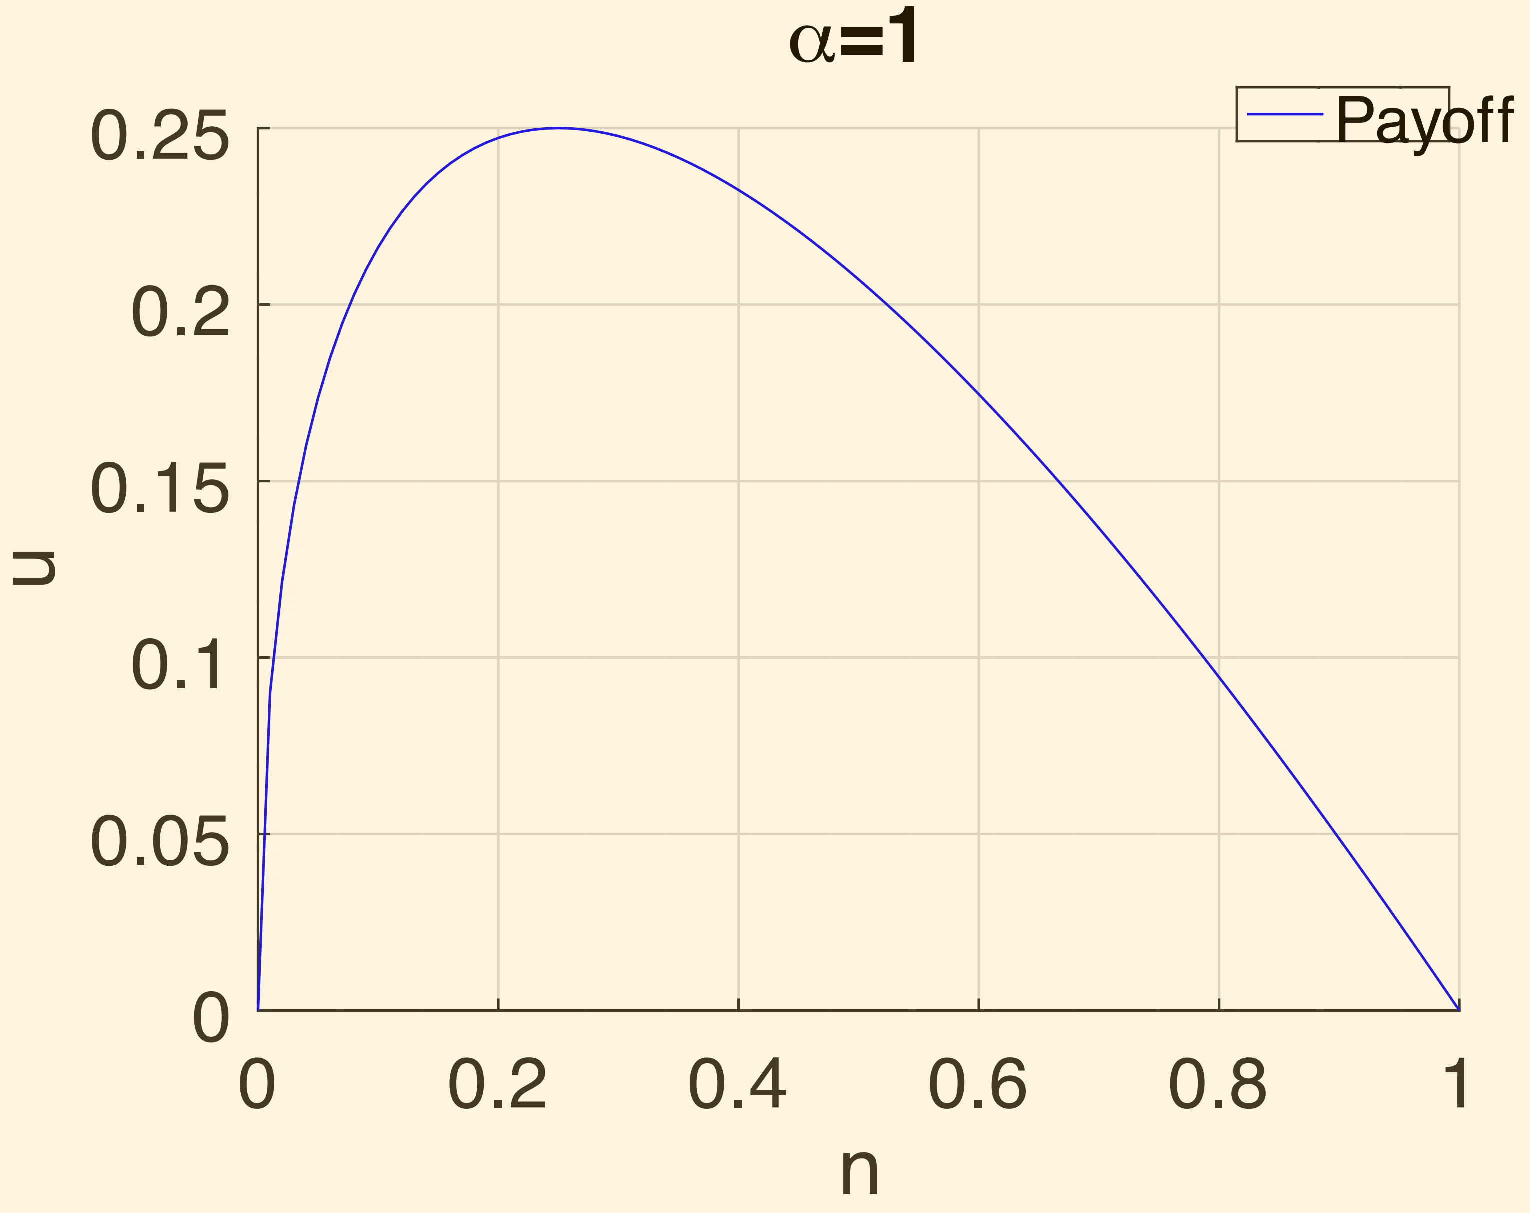

Tragedy of the Data Scientists

- But each of us could have done much better!
- Let's each collect $n^\dagger$ points

$$u_i(n_i = n^\dagger, n_{-i} = n^\dagger) = \sqrt{mn^\dagger} - \alpha n^\dagger$$

$$\frac{du_i(n_i = n^\dagger, n_{-i} = n^\dagger)}{dn^\dagger} = 0$$

$$n^\dagger = \frac{m}{4\alpha^2}$$

Tragedy of the Data Scientists

- $n^\dagger = \dfrac{m}{4\alpha^2}$ is the socially optimal fair assignment (SOFA), $u_i(n^\dagger \ldots n^\dagger) = \dfrac{m}{4\alpha}$

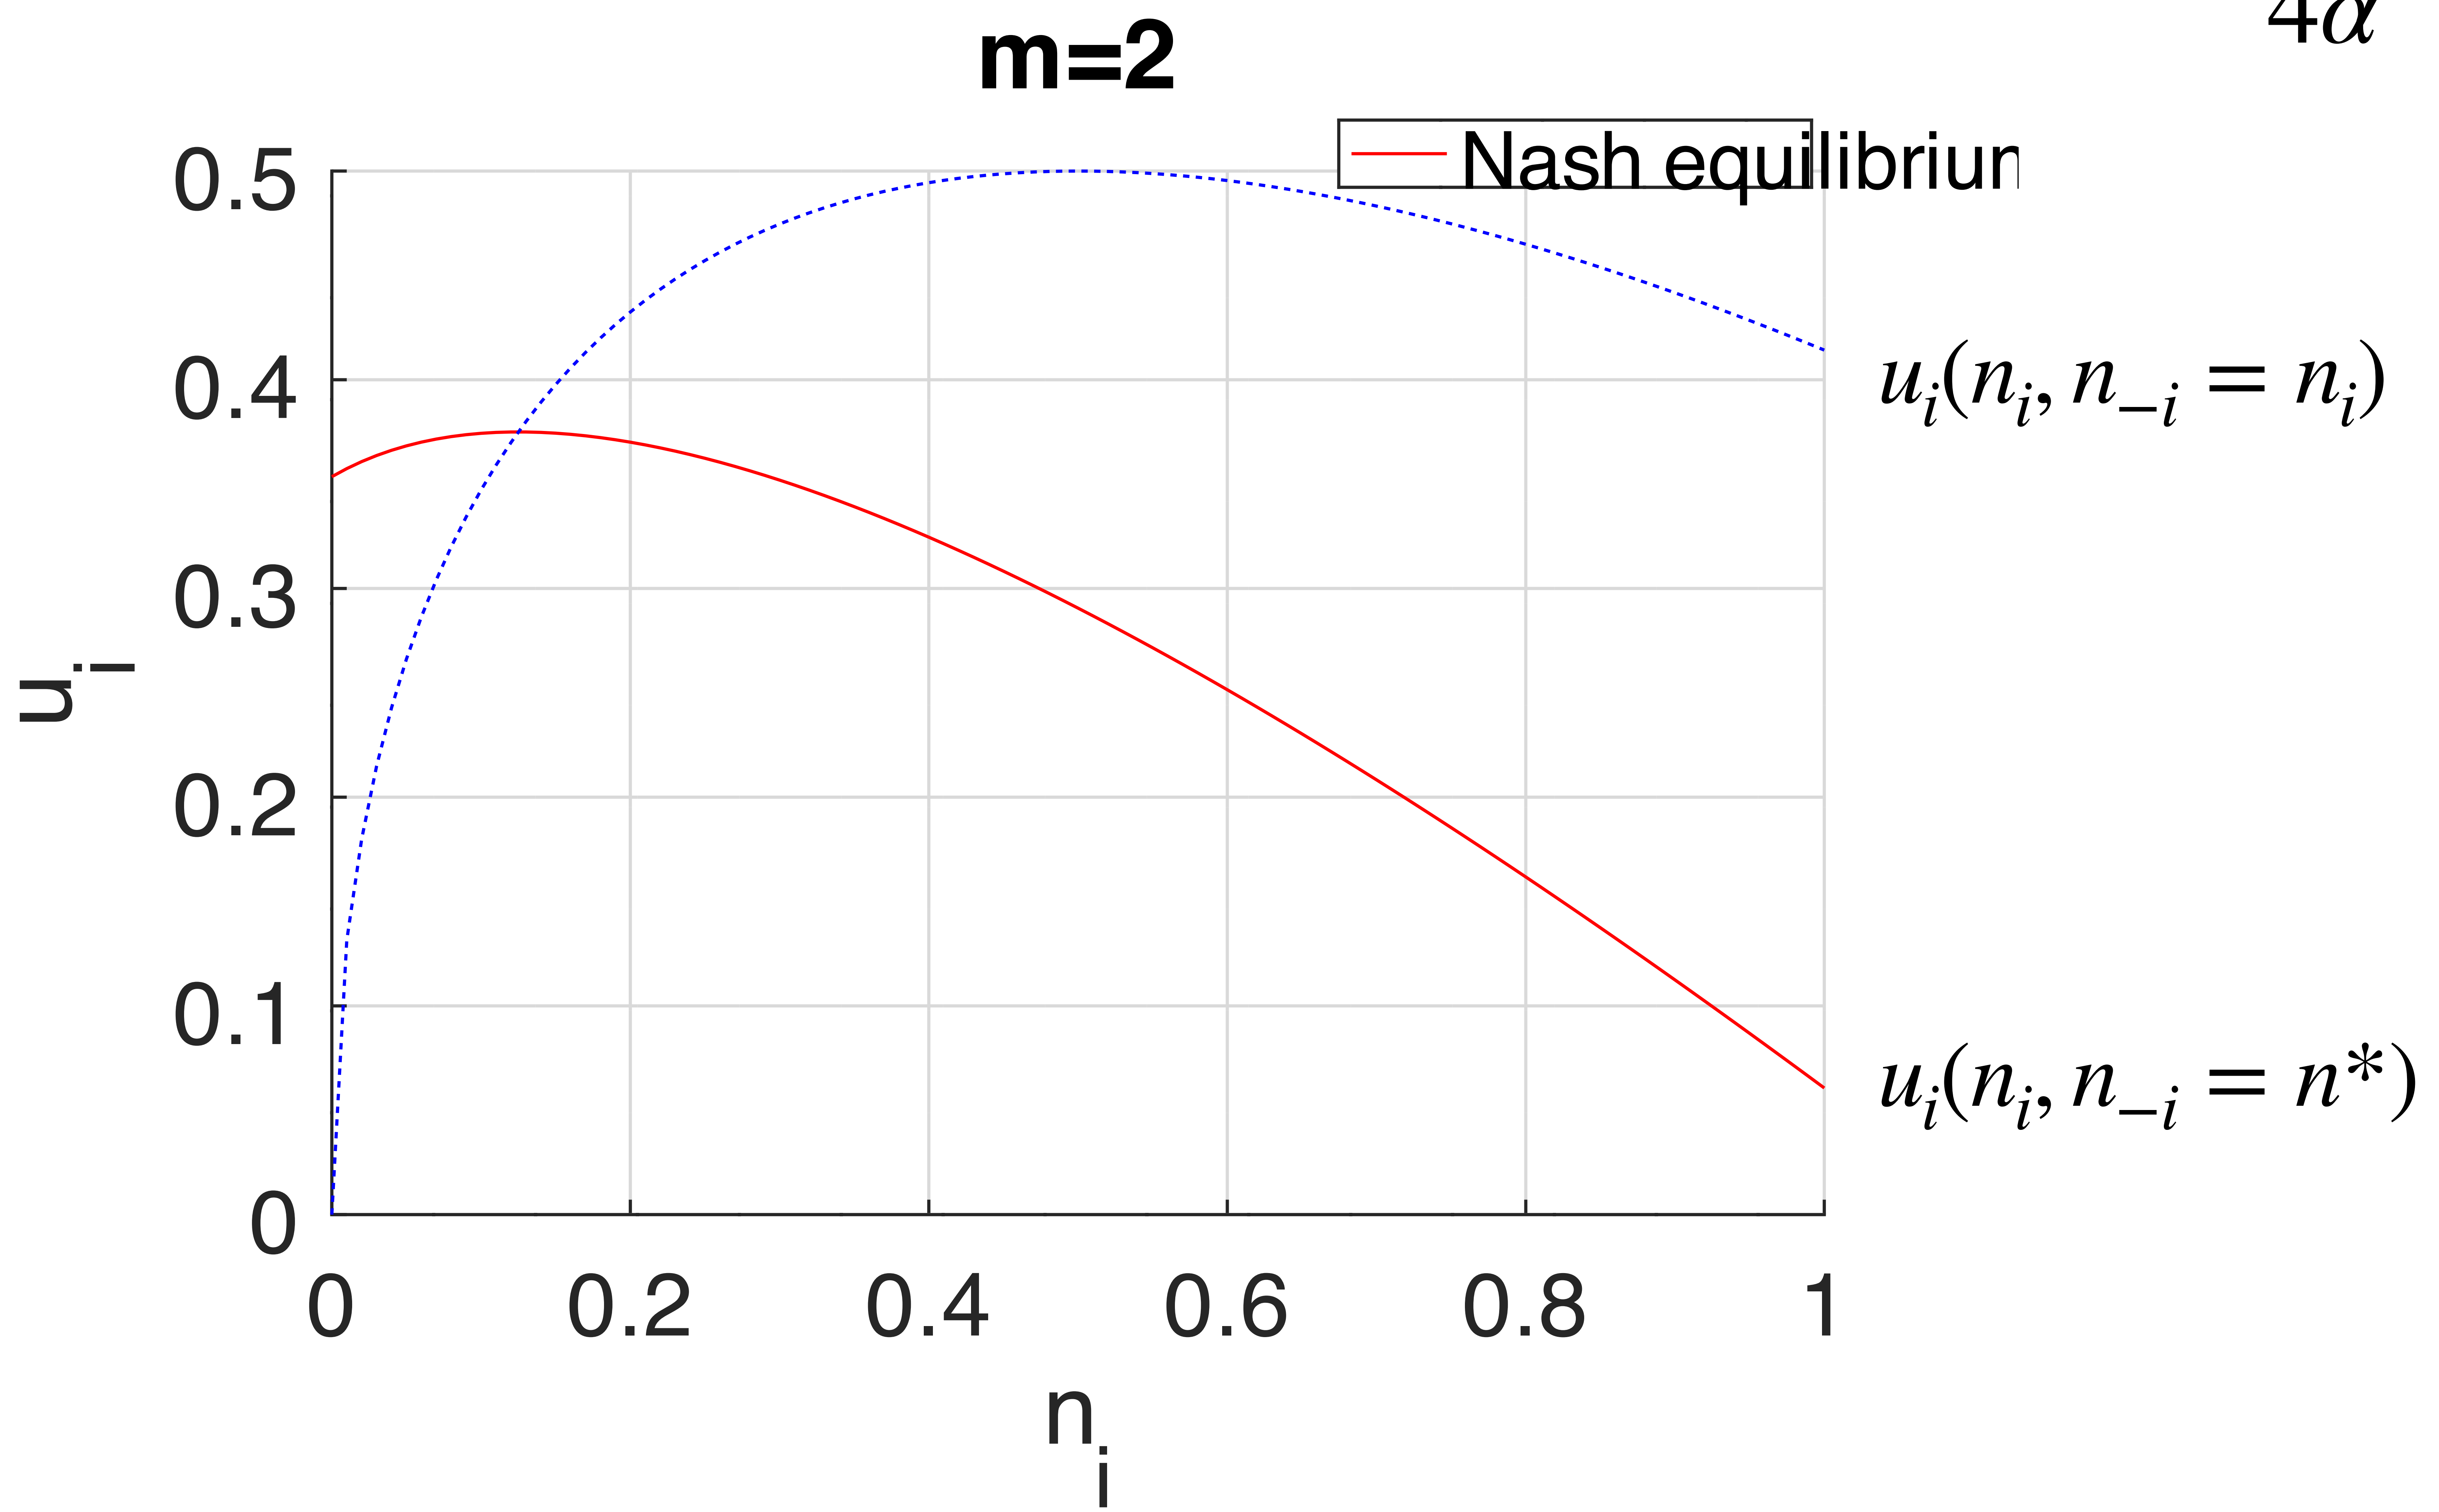

Tragedy of the Data Scientists

- The tragedy: no one will play SOFA.
- If others collect $n^\dagger$, I want to cheat and collect *no data*

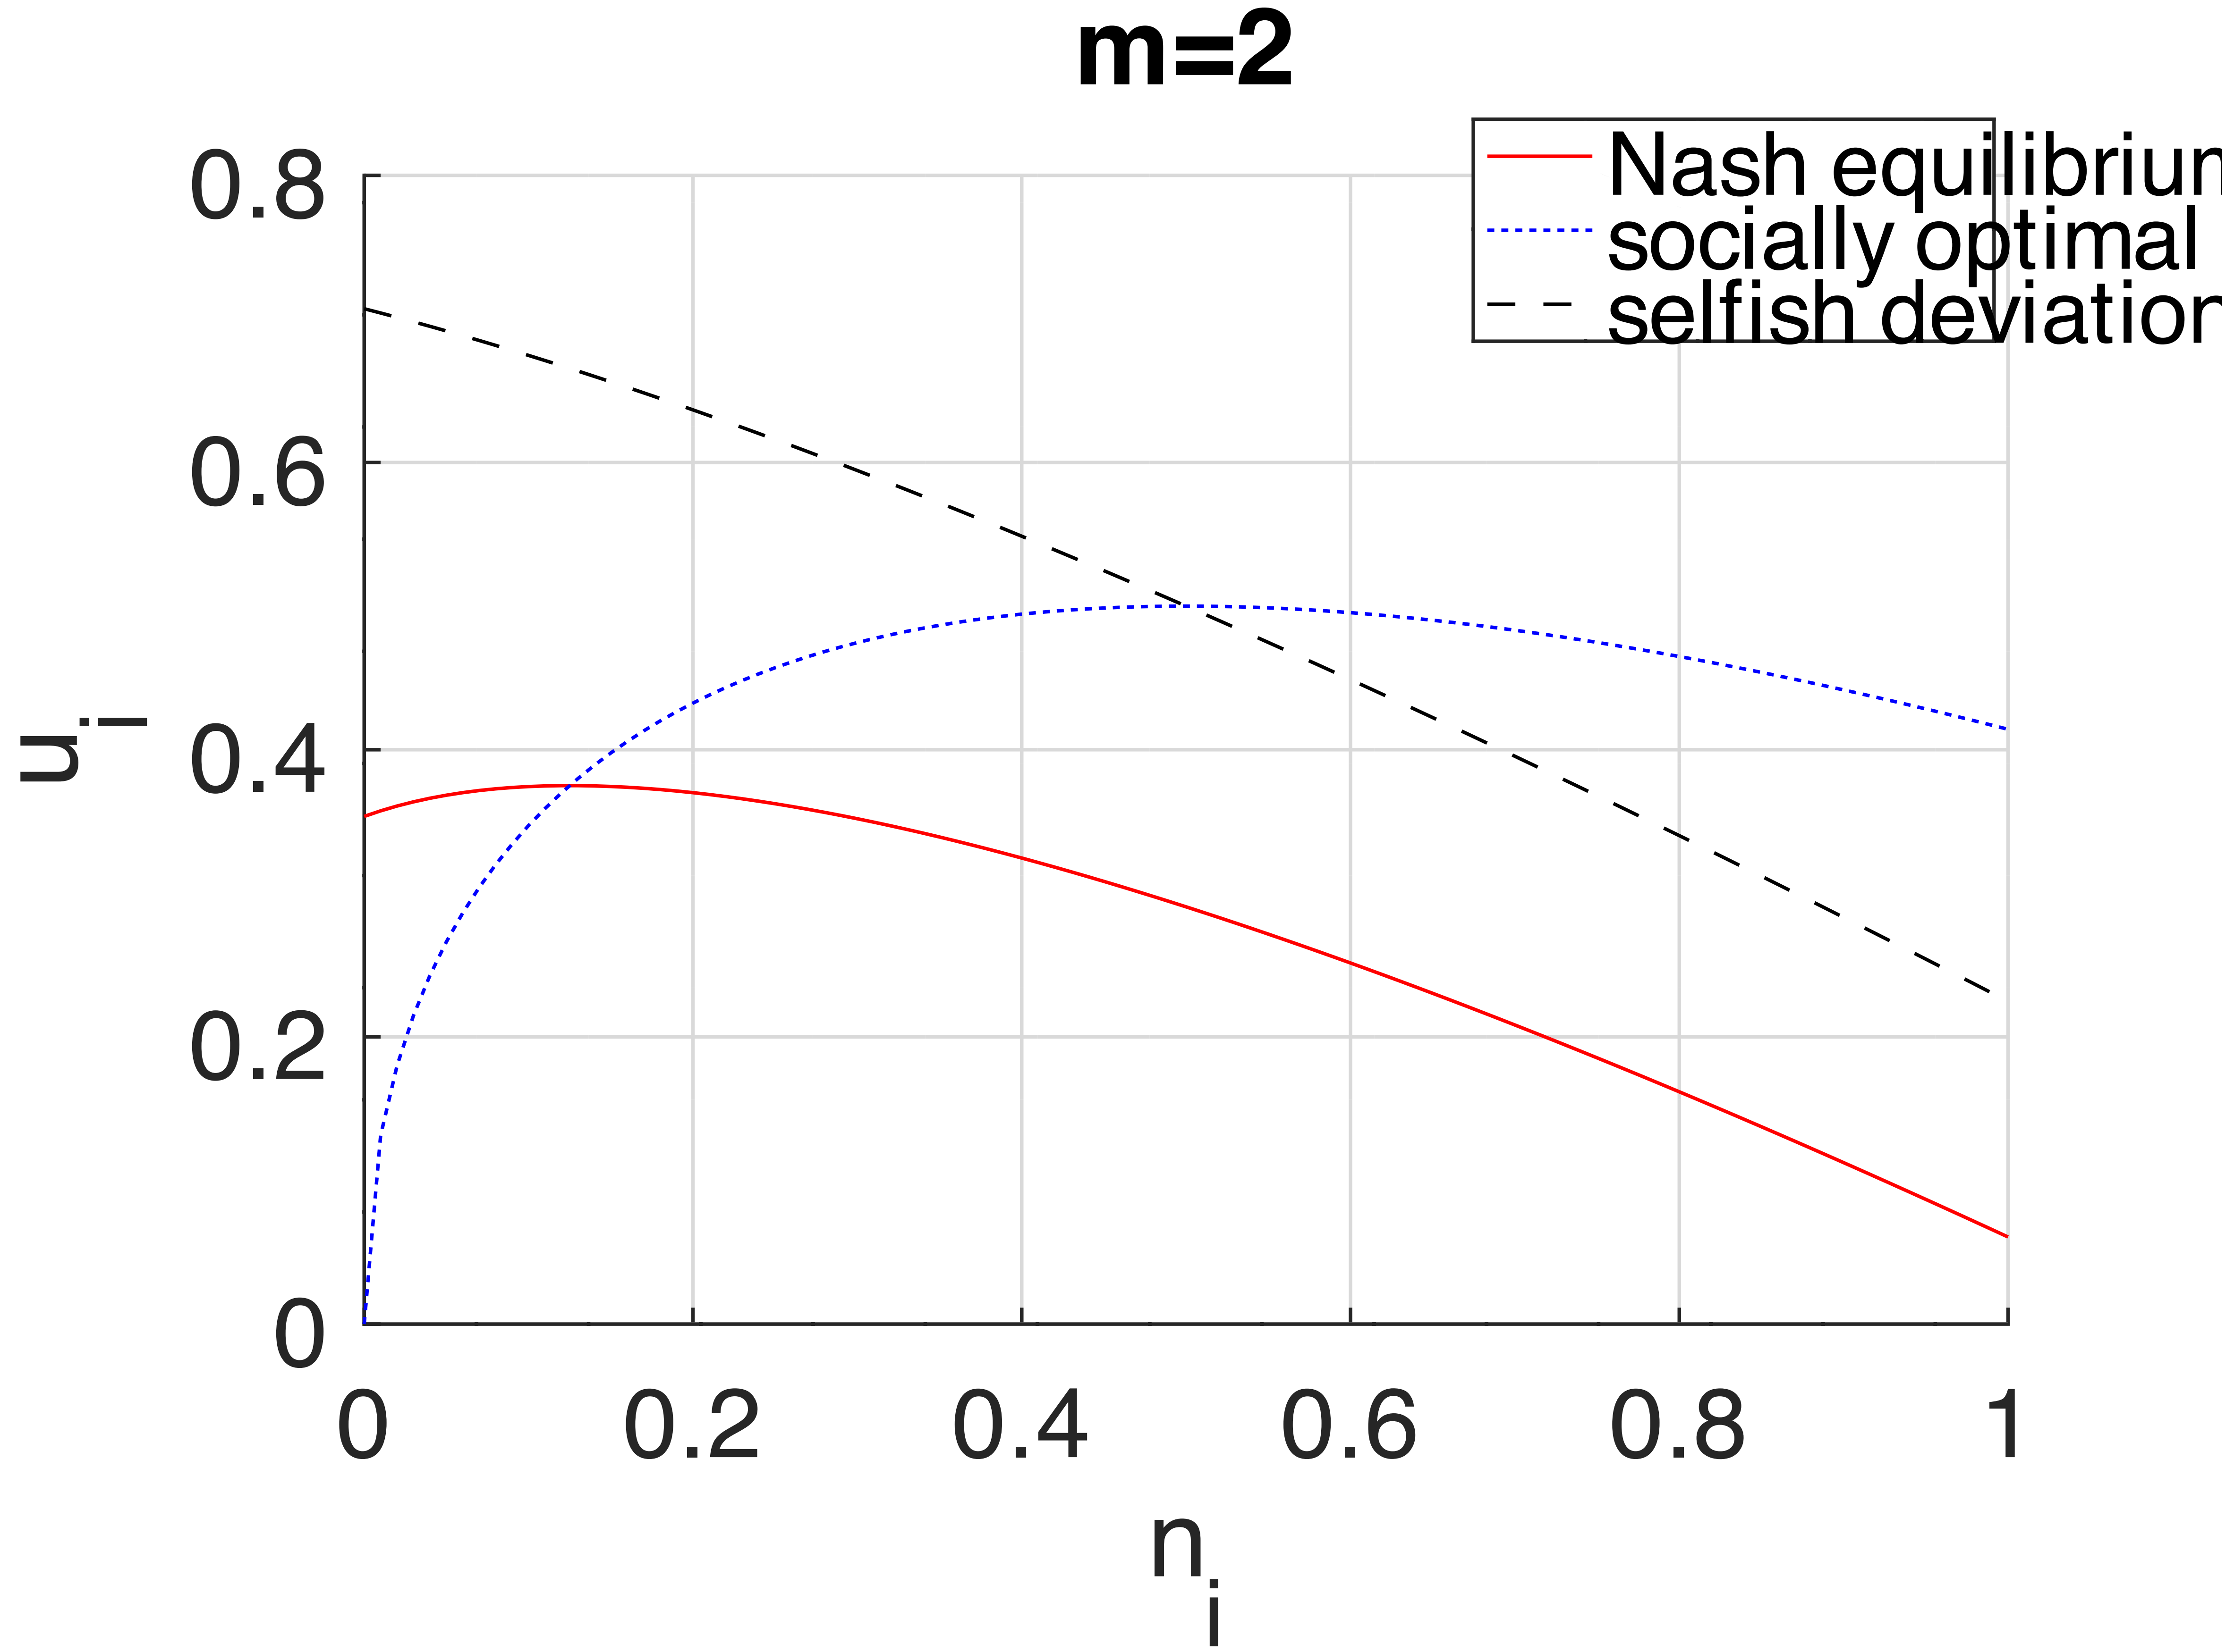

$u_i(n_i, n_{-i} = n_i)$

$u_i(n_i, n_{-i} = n^\dagger)$

$u_i(n_i, n_{-i} = n^*)$

More than Fish

Another example: multi-armed bandit

- Two Gaussian arms $\Delta = \mu_2 - \mu_1 > 0$
- m agents running the ETC algorithm
- Agent j pulls arm 1 $T_j/2$ time and arm 2 $T_j/2$ times (each pull costs α)
- Agents pool data together to find the best empirical arm $\hat{a} \in \{1,2\}$
- Each agent commits to $\hat{a}$ for T' deployment rounds

$$u_j(T_1, \ldots, T_m) = \frac{T_j}{2}(\mu_1 - \alpha) + \frac{T_j}{2}(\mu_2 - \alpha) + T'(E[\mu_{\hat{a}}] - \alpha)$$

More than Fish

Another example: multi-armed bandit

- Nash equilibrium

$$T^* = \frac{4}{\Delta^2 m} W\left(\frac{T'^2 \Delta^6}{32\pi(2\alpha - \mu_1 - \mu_2)^2}\right)$$

- SOFA

$$T^\dagger = \frac{4}{\Delta^2 m} W\left(\frac{m^2 T'^2 \Delta^6}{32\pi(2\alpha - \mu_1 - \mu_2)^2}\right)$$

Lambert W function: $xe^x = z \Rightarrow x = W(z)$

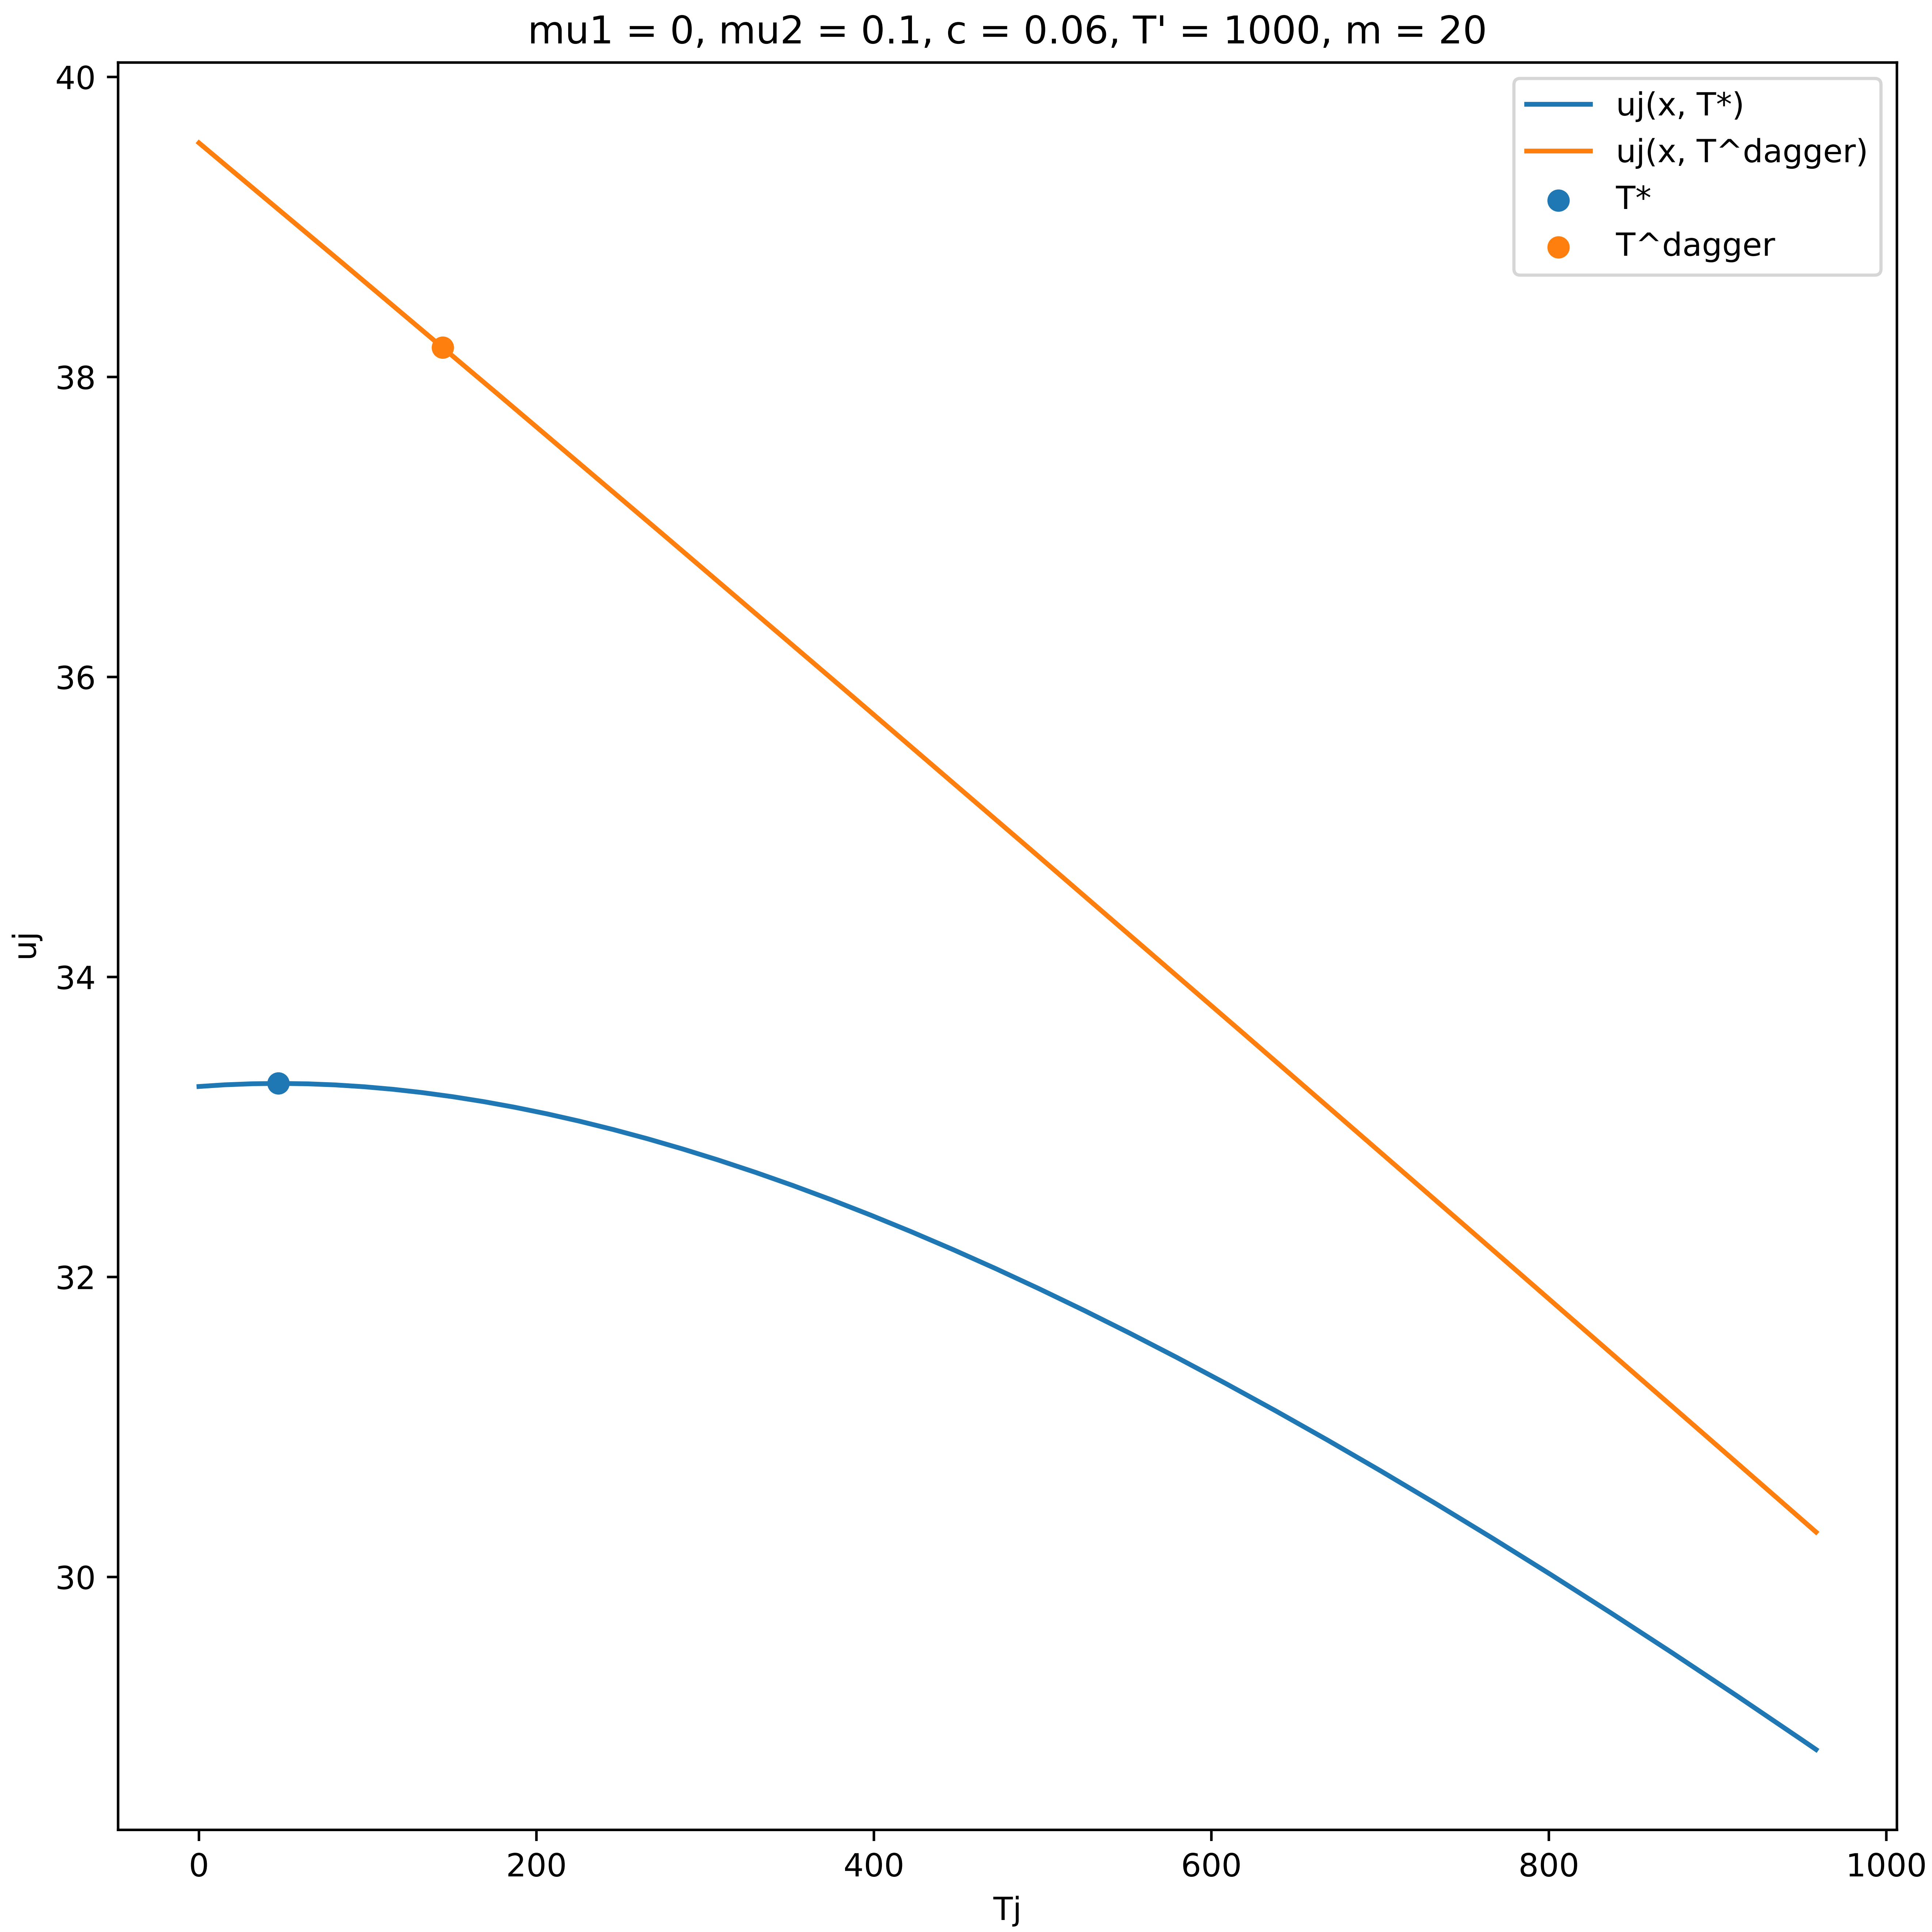

$$u_i(T_i, T_{-i} = T^{\dagger})$$

$$u_i(T_i, T_{-i} = T^{*})$$

Data Collection Inefficiency in General

- Recall payoff = benefit - cost

$$u_i(n_1, \ldots, n_m) := b\left(\sum_{j=1}^{m} n_j\right) - c(n_i)$$

- Nash is the solution to $b'(mn^*) - c'(n^*) = 0$
- SOFA is the solution to $mb'(mn^\dagger) - c'(n^\dagger) = 0$

A Sufficient Condition for Tragedy

$$u_i(n_1, \ldots, n_m) := b\left(\sum_{j=1}^{m} n_j\right) - c(n_i)$$

If:

- b strictly concave and non-decreasing (diminishing return)
- $c = \alpha n_i$ linear (unit cost)

then

$$n^* < n^\dagger$$

Proof by Picture

- Nash: $b'(mn^*) = \alpha$
- SOFA: $b'(mn^\dagger) = \dfrac{\alpha}{m}$

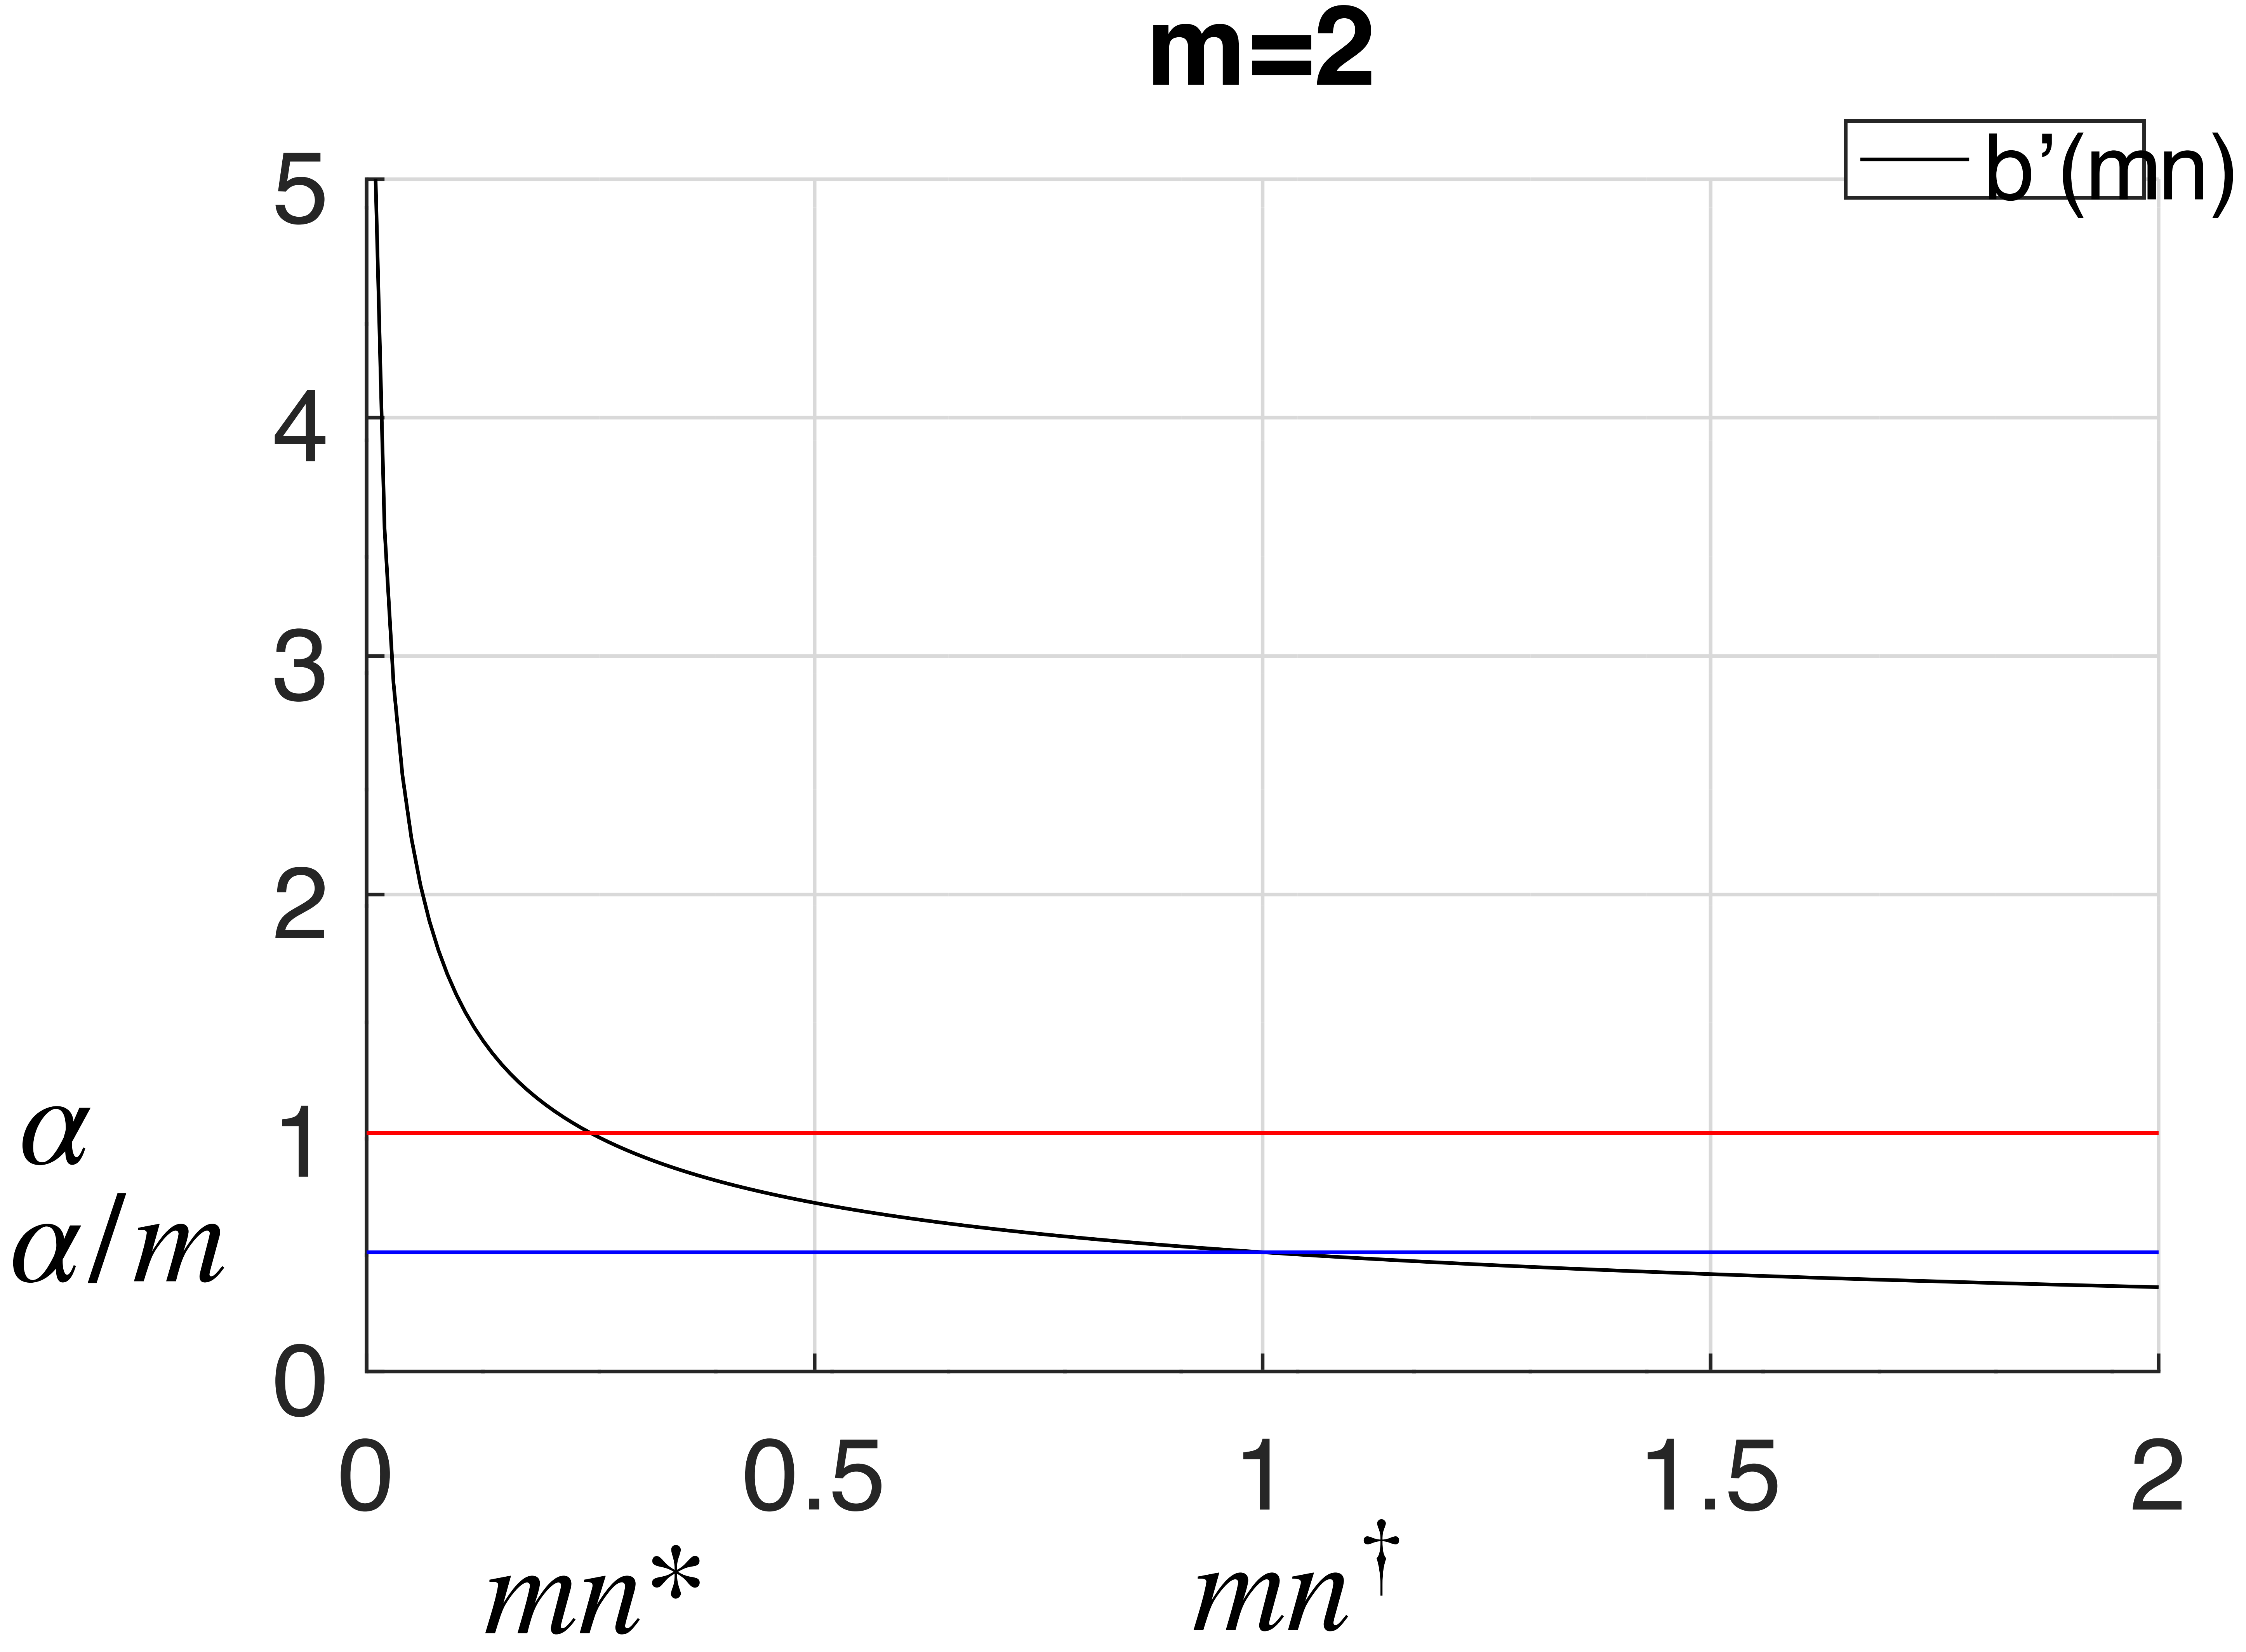

Efficient Data Collection Requires ~~Incentives~~ Coercion

- Paying them a unit price p is not a solution: merely changes α into $\alpha - p$
- Server: takes all data, learn, sends model to agent i only if $n_i = n^\dagger$
  - Changes payoff $u_i(n_1, \ldots, n_m) = b(\sum n_j)\mathbf{1}[n_i = n^\dagger] - c(n_i)$
  - $(n^\dagger, \ldots, n^\dagger)$ now a Nash equilibrium
- Tyranny of the server: enslaves the agents by pushing $n^\dagger$ toward the solution to $u_i = b(mn^\dagger) - c(n^\dagger) \downarrow 0$
- What if agents fake data?