Announcements

- Project proposal and HW2 due today

- Next Friday at Northwestern: Challenges in Data Economics Workshop
  - Speakers: Jon Kleinberg (Cornell), Alessandro Bonatti (MIT), Erik Madsen (NYU), Rachel Cummings (Columbia)
  - Welcome to attend if interested

CMSC 35401: The Interplay of Learning and Game Theory (Autumn 2022)

Crowdsourcing Information and Peer Prediction

Instructor: Haifeng Xu

Outline

- Eliciting Information without Verification
- Equilibrium Concept and Peer Prediction Mechanism
- Bayesian Truth Serum

Crowdsourcing Information

- ➢Recruit AMT workers to label images
  - Cannot check ground truth (too costly)

amazon
beta
mechanical turk

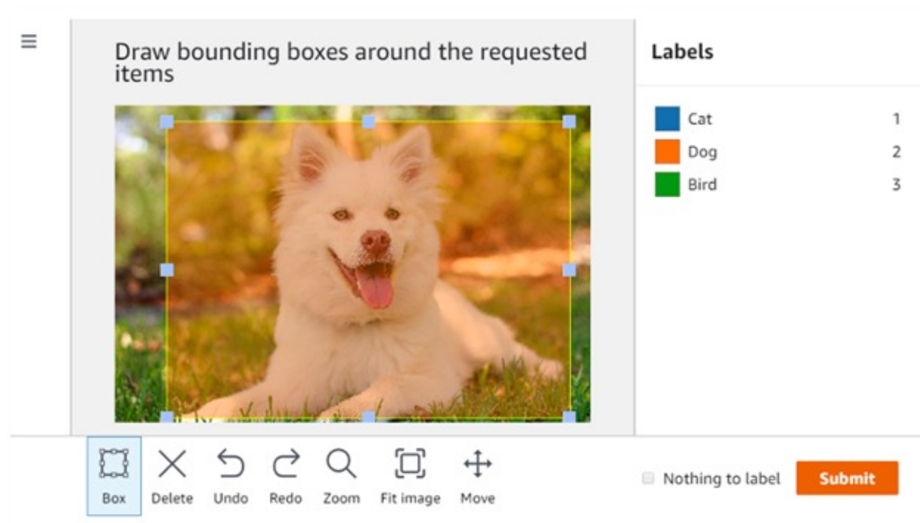

Crowdsourcing Information

- Recruit AMT workers to label images
  - Cannot check ground truth (too costly)

- Peer grading (of, e.g., essays) on MOOC
  - Don't know true scores

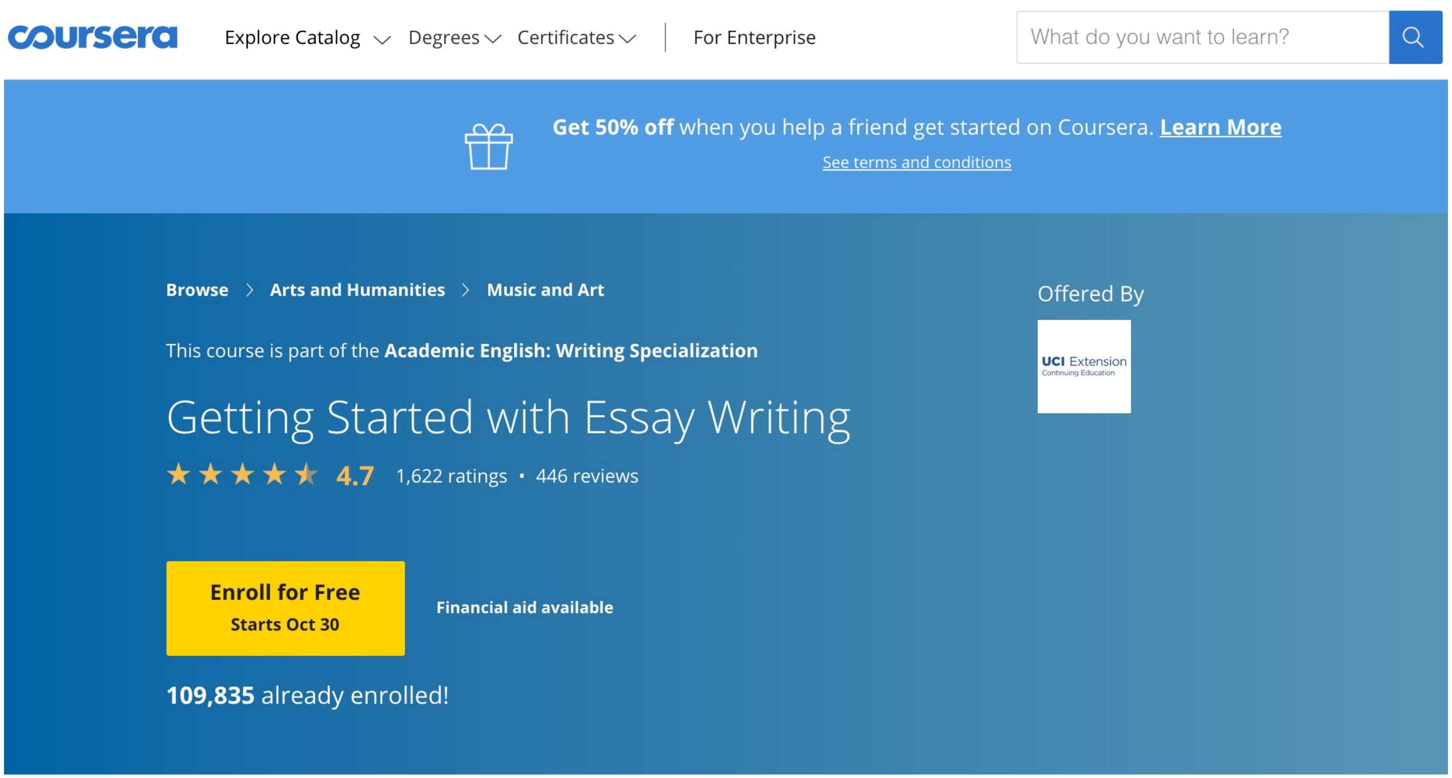

Crowdsourcing Information

- ➢Recruit AMT workers to label images
  - Cannot check ground truth (too costly)

- ➢Peer grading (of, e.g., essays) on MOOC
  - Don’t know true scores

- ➢Elicit ratings for various entities (e.g., on Yelp or Google)
  - We never find out the true quality/rating

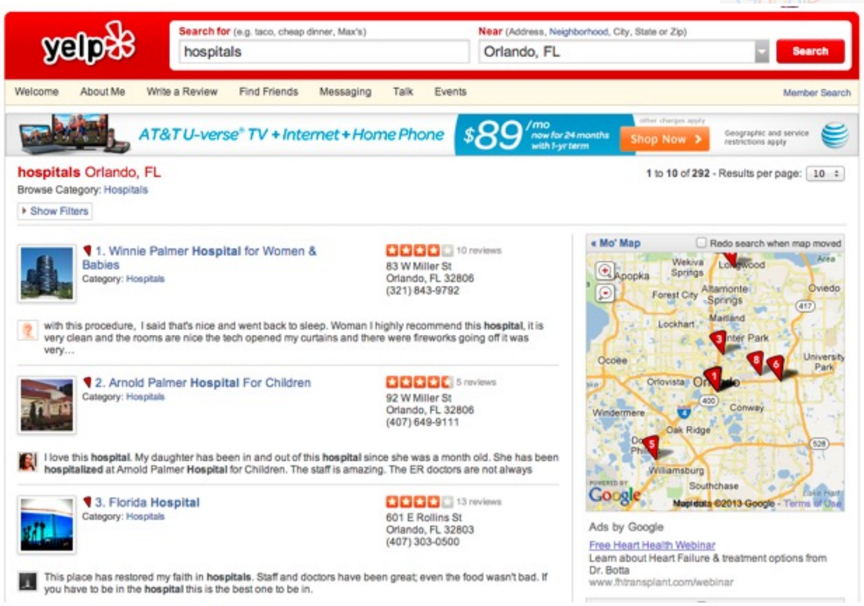

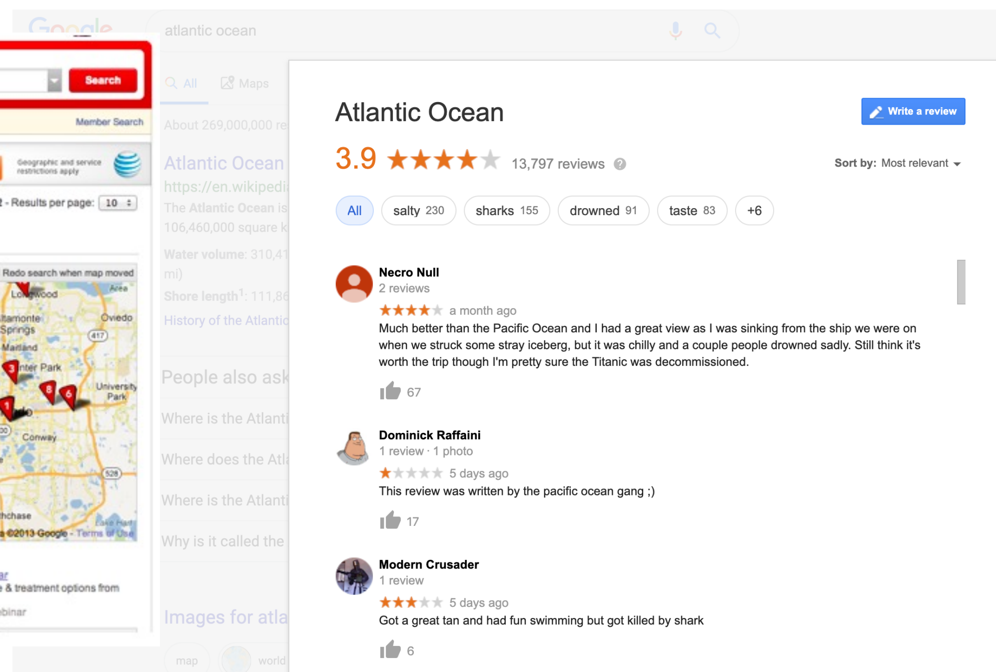

Crowdsourcing Information

- Recruit AMT workers to label images
  - Cannot check ground truth (too costly)

- Peer grading (of, e.g., essays) on MOOC
  - Don’t know true scores

- Elicit ratings for various entities (e.g., on Yelp or Google)
  - We never find out the true quality/rating

- And many other applications…

Common Characteristics in These Applications

- We (the designer) elicit information from population
- Cannot or too costly to know ground truth
  - The reason of using crowdsourcing info elicitation
  - Key difference from prediction markets
- Agents/experts may misreport

Challenge: cannot verify the report/prediction

Solution: let multiple agents compete for the same task, and score them against each other (thus the name "peer prediction")

Another place you see this idea is *auction design*

A Simple and Concrete Example

- Elicit Alice's and Bob's truthful rating A, B about UC dinning
  - $A, B \in \{High, Low\}$
  - There is a common joint belief: $P([A,B] = [H,H]) = 0.5$; $P([A,B] = [H,L]) = 0.24$; $P([A,B] = [L,H]) = 0.24$; $P([A,B] = [L,L]) = 0.02$

Let's try to understand this distribution …

- It is symmetric among Alice and Bob
- $P(A = H) = 0.5 + 0.24 = 0.74$
  - Each expert very likely rates H
- $P(A = H | B = H) = \frac{P(A=H,B=H)}{P(B=H)} = \frac{0.5}{0.74} = \frac{25}{37}$
  - Given that one rates H, the other very likely rates H as well
- $P(A = H | B = L) = \frac{P(A=H,B=L)}{P(B=L)} = \frac{0.24}{0.26} = \frac{12}{13}$
  - Given that one rates L, the other still very likely rates H

A Simple and Concrete Example

- Elicit Alice's and Bob's truthful rating A, B about UC dinning
  - $A, B \in \{High, Low\}$
  - There is a common joint belief: $P([A,B] = [H,H]) = 0.5$; $P([A,B] = [H,L]) = 0.24$; $P([A,B] = [L,H]) = 0.24$; $P([A,B] = [L,L]) = 0.02$
  - $P(A = H) = 0.74$; $P(A = H | B = H) = \frac{25}{37}$; $P(A = H | B = L) = \frac{12}{13}$

Q: What are some natural peer comparison and rewarding mechanisms?

- One natural idea is to reward agreement
  - Ask Alice and Bob to report their signals $\bar{A}$, $\bar{B}$ (may misreport)
  - Award 1 to both if $\bar{A} = \bar{B}$, otherwise reward 0
- Does this work?
  - If $A = H$, what should Alice report?
  - If $A = L$, what should Alice report?

Truthful report is not an equilibrium!

A Simple and Concrete Example

- Elicit Alice's and Bob's truthful rating A, B about UC dinning
  - $A, B \in \{High, Low\}$
  - There is a common joint belief: $P([A, B] = [H, H]) = 0.5$; $P([A, B] = [H, L]) = 0.24$; $P([A, B] = [L, H]) = 0.24$; $P([A, B] = [L, L]) = 0.02$
  - $P(A = H) = 0.74$; $P(A = H | B = H) = \frac{25}{37}$; $P(A = H | B = L) = \frac{12}{13}$

Q: What are some natural peer comparison and rewarding mechanisms?

- Both players always report H (i.e., $\bar{A} = \bar{B} = H$) is a Nash Equ.
- Why?
  - Well, under "rewarding agreement", they both get 1, the maximum possible
  - In fact, both always reporting L is also a NE

Outline

- Eliciting Information without Verification
- Equilibrium Concept and Peer Prediction Mechanism
- Bayesian Truth Serum

The Model of Peer Prediction

- Two experts Alice and Bob, each holding a signal $A \in \{A_1, \cdots, A_n\}$ and $B \in \{B_1, \cdots, B_m\}$ respectively
  - A joint distribution p of (A, B) is publicly known
  - Everything we describe generalize to n experts
- We would like to elicit Alice's and Bob's true signals
  - Despite we never know what signals they truly have

A seemingly richer but equivalent model

- We want to estimate distribution of random var E
- Joint prior distribution p of (A, B, E) is publicly known
  - E.g., E is true quality of our dinning, which we never observe
- Goal: elicit A, B to "refine" our estimation of E (as $E|A, B$)

A Subtle Issue

A seemingly richer but equivalent model

- We want to estimate distribution of random var E
- Joint prior distribution p of (A, B, E) is publicly known
  - E.g., E is true quality of our dinning, which we never observe
- Goal: elicit A, B to refine our estimation of E (as $E|A, B$)

Eliciting signals vs distributions

- In prediction markets, we asked experts to report distributions
- Here, could have done the same thing
  - Alice could report $p(E|A)$, the dist. of E conditioned on her signal A
  - Let's make a minor assumption: $p(E|A) \neq p(E|A')$ for any $A \neq A'$
  - Then, reporting signal A is equivalent to reporting distribution $p(E|A)$
  - So, w.l.o.g., eliciting signals is equivalent
- Drawback: have to assume an accurate and known prior

Info Elicitation Mechanisms and Equilibrium

- Recall, we elicit info by asking Alice's and Bob's signal $\bar{A}$, $\bar{B}$
- As before, will design rewards $r_A(\bar{A}, \bar{B})$ and $r_B(\bar{A}, \bar{B})$
- Alice's action is a report strategy $\sigma_A(A) \in \{A_1, \cdots, A_n\}$ [Bob similar]
  - This is a pure strategy
  - Will not consider mixed strategy here, because we will design r_A and r_B so that there is a good pure equilibrium
  - Truth-telling strategy: $\sigma_A(A) = A, \sigma_B(B) = B$
- Then, what outcome is expected to occur? → equilibrium outcome
- Generally, it is a Bayesian Nash equilibrium (**BNE**)
  - For simplicity, only define BNE here for our particular setting

Info Elicitation Mechanisms and Equilibrium

- Recall, we elicit info by asking Alice's and Bob's signal $\bar{A}, \bar{B}$
- As before, will design rewards $r_A(\bar{A}, \bar{B})$ and $r_B(\bar{A}, \bar{B})$
- Alice's action is a report strategy $\sigma_A(A) \in \{A_1, \cdots, A_n\}$ [Bob similar]

Definition. $\sigma_A(A), \sigma_B(B)$ is a Bayesian Nash equilibrium if the following holds

$$\forall A \qquad \mathbb{E}_{B|A}\, r_A\big(\sigma_A(A), \sigma_B(B)\big) \geq \mathbb{E}_{B|A}\, r_A\big(\sigma'_A(A), \sigma_B(B)\big),$$

$$\forall B \qquad \mathbb{E}_{A|B}\, r_B\big(\sigma_A(A), \sigma_B(B)\big) \geq \mathbb{E}_{A|B}\, r_B\big(\sigma_A(A), \sigma'_B(B)\big).$$

We say it is a strict BNE if both "≥" are ">"

Mechanism for Peer Prediction

- Design objective: choose r_A, r_B so that truth-telling is an Equ.

Any ideas?

- Use proper scoring rules? But don't have a realized outcome…
- Key idea: Alice's signal can be used to score the distribution induced by Bob's signal, and vice versa

Mechanism for Peer Prediction

Information Elicitation without Verification

"Parameter": any strict proper scoring rule $S(i; p)$

1. Elicit Alice's signal $\bar{A}$ and Bob's signal $\bar{B}$
2. Calculate $p_{\bar{A}}(B) = \text{dist. of } B \text{ conditioned on } \bar{A}$, and similarly $p_{\bar{B}}(A)$
3. Award Alice $r_A(\bar{A}, \bar{B}) = S(\bar{B}; p_{\bar{A}})$ and Bob $r_B(\bar{A}, \bar{B}) = S(\bar{A}; p_{\bar{B}})$

Remark:

- Step 2 relies on the prior distribution
- Alice is awarded by how accurate her report $\bar{A}$ (equivalently $p_{\bar{A}}$) predicts Bob's B.

Mechanism for Peer Prediction

Information Elicitation without Verification

"Parameter": any strict proper scoring rule $S(i; p)$

1. Elicit Alice's signal $\bar{A}$ and Bob's signal $\bar{B}$
2. Calculate $p_{\bar{A}}(B) = \text{dist. of } B \text{ conditioned on } \bar{A}$, and similarly $p_{\bar{B}}(A)$
3. Award Alice $r_A(\bar{A}, \bar{B}) = S(\bar{B}; p_{\bar{A}})$ and Bob $r_B(\bar{A}, \bar{B}) = S(\bar{A}; p_{\bar{B}})$

Theorem. Truth-telling is a strict BNE in the above game

Proof: show $\sigma_A(A) = A$ is a best response to $\sigma_B(B) = B$, and vice versa

- If Bob reports B truthfully, Alice receives $S(B; p_{\bar{A}})$ by reporting $\bar{A}$
- With true signal A, what is Alice's best response report $\bar{A}$?
    - By strict properness, Alice wants $p_{\bar{A}}$ to be exactly her true belief of dist. of B
    - So, Alice should report $\bar{A} = A$.

Remarks

- Mechanism is only described for two experts, but no difficult to generalize to n experts
  - Can randomly match each expert to a "peer" as reference
- Serious issues are the following

Issue 1: there are many other equilibria in the game

- Dinning rating example with slightly different numbers
  - A common joint belief: $P([A, B] = [H, H]) = 0.4$; $P([A, B] = [H, L]) = 0.1$; $P([A, B] = [L, H]) = 0.1$; $P([A, B] = [L, L]) = 0.4$
- Both always report H is also an equilibrium
  - If Bob reports H, Alice's reward is $S(H; p_{\bar{A}})$ (regardless of her true A)
  - $\bar{A} = H$ makes $p_{\bar{A}}(H) = P(B = H | \bar{A} = H) = 4/5$
  - $\bar{A} = L$ makes $p_{\bar{A}}(H) = P(B = H | \bar{A} = L) = 1/5$

Remarks

- Mechanism is only described for two experts, but no difficult to generalize to n experts
  - Can randomly match each expert to a "peer" as reference
- Serious issues are the following

Issue 1: there are many other equilibria in the game

- More generally, reporting signals that are easy to coordinate likely forms an equilibrium
  - E.g., you are asked to grade essays, but you may all report the length of the essay while not its true quality (less effort, more well correlated)
- This is a fundamental issue of peer prediction

Key challenge: how to design mechanisms where truth-telling is unique (or the most plausible) equilibrium

Remarks

- Mechanism is only described for two experts, but no difficult to generalize to n experts
  - Can randomly match each expert to a "peer" as reference
- Serious issues are the following

Issue 2: Designer has to know the joint distribution of (A, B)

- Not very realistic, as designer usually has little knowledge
- But, there are remedies for this issue…

Outline

- Eliciting Information without Verification
- Equilibrium Concept and Peer Prediction Mechanism
- Bayesian Truth Serum

Designed for a Special yet Realistic Setting

- We, the designer, want to predict distribution of E
- n experts, each i has a signal $S_i \sim p(S|E)$ i.i.d.
  - In this setting, we have to have many experts
  - Assume experts know $p(S|E)$ but we do not know
- Objective: elicit true signals $S_1, \cdots, S_n$

Key design ideas

Designed for a Special yet Realistic Setting

- We, the designer, want to predict distribution of E
- n experts, each i has a signal $S_i \sim p(S|E)$ i.i.d.
  - In this setting, we have to have many experts
  - Assume experts know $p(S|E)$ but we do not know
- Objective: elicit true signals $S_1, \cdots, S_n$

Key design ideas

- Cannot compute posterior distribution conditioned on any expert's signal anymore, but still need it to score him
- So, will elicit both his signal and his posterior belief of others' signals

Bayesian Truth Serum [Prelec, Science'04]

The Protocol

1. For each i, elicit her signal $\bar{S}_i$ and her prediction $\bar{p}^i \in \Delta_{|S|}$ of the distribution of other expert's signal (agents are i.i.d. a-priori)
2. Calculate (geometric) mean prediction $\bar{p}$ where
$$\bar{p}_S = \sqrt[n]{\bar{p}_S^1 \times \bar{p}_S^2 \cdots \bar{p}_S^n} \qquad \text{for any signal } S$$
3. Compute $\bar{\lambda}$ as the empirical distribution of reported signals $\bar{S}_i$'s.
4. Reward agent i the following (G is any proper scoring rule)
$$\log \frac{\bar{\lambda}_{\bar{S}_i}}{\bar{p}_{\bar{S}_i}} + \boxed{\mathbb{E}_{S\sim\bar{\lambda}}\, G(S; \bar{p}^i)}$$

Score of i's prediction $\bar{p}^i$, against the true signal distribution $\bar{\lambda}$

- By properness, want $\bar{p}^i$ to be close to $\bar{\lambda}$

Bayesian Truth Serum [Prelec, Science'04]

The Protocol

1. For each i, elicit her signal $\bar{S}_i$ and her prediction $\bar{p}^i \in \Delta_{|S|}$ of the distribution of other expert's signal (agents are i.i.d. a-priori)
2. Calculate (geometric) mean prediction $\bar{p}$ where
   $$\bar{p}_S = \sqrt[n]{\bar{p}_S^1 \times \bar{p}_S^2 \cdots \bar{p}_S^n} \qquad \text{for any signal } S$$
3. Compute $\bar{\lambda}$ as the empirical distribution of reported signals $\bar{S}_i$'s.
4. Reward agent i the following (G is any proper scoring rule)
   $$\log \frac{\bar{\lambda}_{\bar{S}_i}}{\bar{p}_{\bar{S}_i}} + \mathbb{E}_{S\sim\bar{\lambda}}\, G(S; \bar{p}^i)$$

Score of i's signal report S_i

➢ This score is large if $\bar{\lambda}_{\bar{S}_i} \geq \bar{p}_{\bar{S}_i}$ --- that is, i's reported type is surprisingly more common than designer's predicted probability $\bar{p}_{\bar{S}_i}$

Bayesian Truth Serum [Prelec, Science'04]

Theorem. When $n \to \infty$, truthful report is a Bayesian Nash equilibrium in the previous protocol.

- That is, expert i should report his true signal S_i and his true posterior belief of other expert's signals
- $n \to \infty$ is needed because in that case $\bar{\lambda} \to$ the exact signal distribution (under truthful signal report)
  - Later works relax this assumption to requiring only large enough n
- Proof is a bit intricate (see the Science paper)
- Very insightful – particularly, the usefulness of rewarding "surprisingly common" signals is not clear before at all
- The issue of existence of multiple equilibria is still there

Thank You

Haifeng Xu
University of Chicago
haifengxu@uchicago.edu