Announcements

- HW and project proposal due this Thur
  - 2pm for HW
  - 6pm for proposal

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

Scoring Rules and its Connection to Prediction Markets

Instructor: Haifeng Xu

Outline

- Scoring Rule and its Characterization
- Connection to Prediction Markets
- Gaming a Prediction Markets

Consider a Simpler Setting

- We (designer) want to learn the distribution of random var $E \in [n]$
  - E will be sampled in the future
- We have no samples from E; Instead, we have an expert/predictor who has a predicted distribution $\lambda \in \Delta_n$
- We want to incentivize the expert to truthfully report λ

Consider a Simpler Setting

- We (designer) want to learn the distribution of random var $E \in [n]$
  - E will be sampled in the future
- We have no samples from E; Instead, we have an expert/predictor who has a predicted distribution $\lambda \in \Delta_n$
- We want to incentivize the expert to truthfully report λ

Example

- E is whether UChicago will win Nobel in 2023
- Expert is a famous Econ prof.

- Expert's prediction does not need to be perfect
  - But, better than the designer who knows nothing
- Assume expert will not give you truthful info for free

Idea: "Score" Expert's Report

Will reward the expert certain amount $S(i; p)$ where:
(1) p is the expert's report (does not have to equal λ);
(2) $i \in [n]$ is the event realization

Setup is not like a prediction market (yet), but will see later they are related

Idea: "Score" Expert's Report

Will reward the expert certain amount $S(i;p)$ where:
(1) p is the expert's report (does not have to equal λ);
(2) $i \in [n]$ is the event realization

Q: what is the expert's expected utility?

- Expert believes $i \sim \lambda$
- Expected utility $\mathbb{E}_{i\sim\lambda} S(i;p) = \sum_{i\in[n]} \lambda_i \cdot S(i;p)$ $:= S(\lambda;p)$

Q: what $S(i;p)$ function can elicit truthful report λ?

- When expert finds that $\lambda = \arg\max_{p\in\Delta_n} [\sum_{i\in[n]} \lambda_i \cdot S(i;p)]$
- Ideally, λ is the unique maximizer

Proper Scoring Rules

Definition. A "scoring rule" $S(i; p)$ is [strictly] proper if truthful report $p = \lambda$ [uniquely] maximizes expected utility $S(\lambda; p)$.

- Expert is incentivized to report truthfully iff $S(i; p)$ is proper

Observations.

1. $S(i; p) = 0$ is a trivial proper scoring fnc
2. Proper scores are closed under affine transformation
   - That is, if $S(i; p)$ is [strictly] proper, so is $\alpha \cdot S(i; p) + \beta$ for any constant $\alpha > 0, \beta$

- Thus, typically, strict properness is desired

Examples of Scoring Rules

Example 1 [Log Scoring Rule]

- $S(i; p) = \log p_i$
- $S(\lambda; p) = \sum_{i \in [n]} \lambda_i \cdot \log p_i$

- Negative, but okay – can always add a constant
- Properness requires $\lambda = \arg\max_{p \in \Delta_n} S(\lambda; p)$ Does this hold?

$$S(\lambda; p) = \sum_{i \in [n]} \lambda_i \cdot \log p_i$$

$$= \sum_{i \in [n]} \lambda_i [\log p_i - \log \lambda_i] + \sum_{i \in [n]} \lambda_i \log \lambda_i$$

Examples of Scoring Rules

Example 1 [Log Scoring Rule]

- $S(i; p) = \log p_i$
- $S(\lambda; p) = \sum_{i \in [n]} \lambda_i \cdot \log p_i$

- Negative, but okay – can always add a constant
- Properness requires $\lambda = \arg\max_{p \in \Delta_n} S(\lambda; p)$ Does this hold?

$$S(\lambda; p) = \sum_{i \in [n]} \lambda_i \cdot \log p_i$$

$$= \sum_{i \in [n]} \lambda_i [\log p_i - \log \lambda_i] + \sum_{i \in [n]} \lambda_i \log \lambda_i$$

$$= -\sum_{i \in [n]} \lambda_i \cdot \log \frac{\lambda_i}{p_i} - Entrop(\lambda)$$

Examples of Scoring Rules

Example 1 [Log Scoring Rule]

- $S(i;p) = \log p_i$
- $S(\lambda;p) = \sum_{i\in[n]} \lambda_i \cdot \log p_i$

- Negative, but okay – can always add a constant
- Properness requires $\lambda = \arg\max_{p\in\Delta_n} S(\lambda;p)$ Does this hold?

$$S(\lambda;p) = \sum_{i\in[n]} \lambda_i \cdot \log p_i$$

$$= \sum_{i\in[n]} \lambda_i [\log p_i - \log \lambda_i] + \sum_{i\in[n]} \lambda_i \log \lambda_i$$

$$= -\underbrace{\sum_{i\in[n]} \lambda_i \cdot \log \frac{\lambda_i}{p_i}} - Entrop(\lambda)$$

KL-divergence $KL(\lambda;p)$ (a.k.a. relative entropy)

- Measures the distance between two distributions
- Always non-negative, and equals 0 only when $p = \lambda$

Examples of Scoring Rules

Example 1 [Log Scoring Rule]

- $S(i; p) = \log p_i$
- $S(\lambda; p) = \sum_{i \in [n]} \lambda_i \cdot \log p_i$

- Negative, but okay – can always add a constant
- Properness requires $\lambda = \arg\max_{p \in \Delta_n} S(\lambda; p)$ Does this hold? ✔

$$S(\lambda; p) = \sum_{i \in [n]} \lambda_i \cdot \log p_i$$

$$= \sum_{i \in [n]} \lambda_i [\log p_i - \log \lambda_i] + \sum_{i \in [n]} \lambda_i \log \lambda_i$$

$$= -\sum_{i \in [n]} \lambda_i \cdot \log \frac{\lambda_i}{p_i} - Entrop(\lambda)$$

- p should minimize distance $KL(\lambda; p)$, which is achieved at $p = \lambda$
- Log scoring rule is strictly proper

Examples of Scoring Rules

Example 2 [Quadratic Scoring Rule]

- $S(i; p) = 2p_i - \sum_{j\in[n]} p_j^2$
- $S(\lambda; p) = \sum_{i\in[n]} \lambda_i [2p_i - \sum_{j\in[n]} p_j^2]$

$$
\begin{aligned}
S(\lambda; p) &= \sum_{i\in[n]} \lambda_i [2p_i - \sum_{j\in[n]} p_j^2] \\
&= \sum_{i\in[n]} 2\lambda_i p_i - \left(\sum_{i\in[n]} \lambda_i\right) \cdot \sum_{j\in[n]} p_j^2 \\
&= \sum_{i\in[n]} 2\lambda_i p_i - \sum_{i\in[n]} p_i^2 \\
&= -\sum_{i\in[n]} [p_i - \lambda_i]^2 + \sum_{i\in[n]} \lambda_i^2
\end{aligned}
$$

- Prediction p should minimize l_2-distance between p and λ
- $p_i = \lambda_i$ is the unique maximizer of $S(\lambda; p)$
- Quadratic scoring rule is also strictly proper

Examples of Scoring Rules

Example 3 [Linear Scoring Rule]

- $S(i; p) = p_i$
- $S(\lambda; p) = \sum_{i \in [n]} \lambda_i p_i$

- Linear scoring rule turns out to be not proper (verify it after class)

What $S(i; p)$ Are Proper?

Theorem. The scoring rule $S(i; p)$ is (strictly) proper if and only if there exists a (strictly) convex function $G: \Delta_n \rightarrow \mathbb{R}$ such that

$$S(i; p) = G(p) + \nabla G(p)(e_i - p)$$

basis vector (0,..,0,1,0,…,0)

Recall $G(p)$ is convex if for any $\alpha \in [0,1]$

$$\alpha G(p) + (1 - \alpha)G(q) \geq G(\ \alpha p + (1 - \alpha)q\ )$$

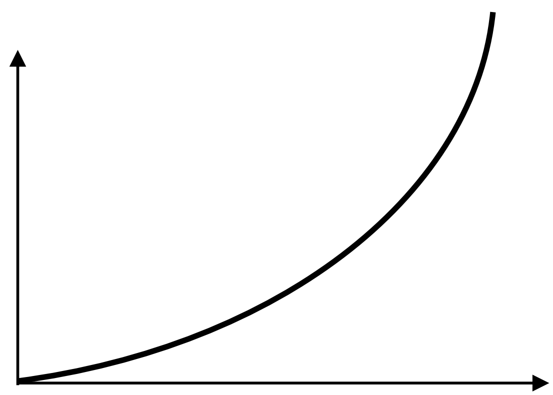

What $S(i;p)$ Are Proper?

Theorem. The scoring rule $S(i;p)$ is (strictly) proper if and only if there exists a (strictly) convex function $G: \Delta_n \to \mathbb{R}$ such that

$$S(i;p) = G(p) + \nabla G(p)(e_i - p)$$

Proof of "⇐"

$$S(\lambda;p) = \mathbb{E}_{i\sim\lambda}[G(p) + \nabla G(p)(e_i - p)]$$

$$= G(p) + \nabla G(p)(\lambda - p)$$

$$\leq G(\lambda) = S(\lambda;\lambda)$$

By convexity

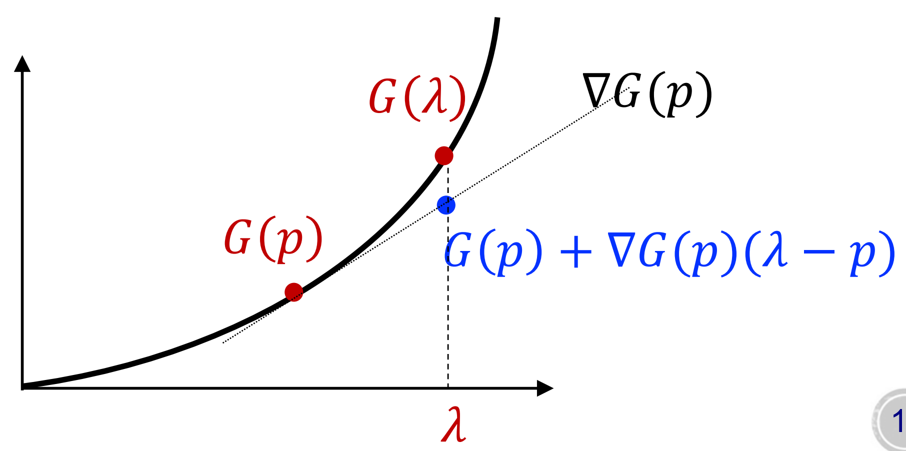

What $S(i;p)$ Are Proper?

Theorem. The scoring rule $S(i;p)$ is (strictly) proper if and only if there exists a (strictly) convex function $G: \Delta_n \to \mathbb{R}$ such that

$$S(i;p) = G(p) + \nabla G(p)(e_i - p)$$

Proof of "⇒"

- $S(\lambda;p) = \sum_{i\in[n]} \lambda_i S(i;p)$ is a linear fnc of λ for any p
- By properness, $S(\lambda;\lambda) = \max_{p\in\Delta_n} S(\lambda;p)$, denoted as $G(\lambda)$
  - $G(\lambda)$ is convex in λ since it is max of linear functions

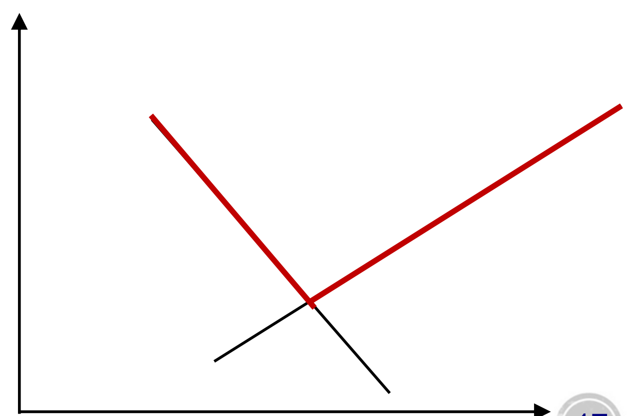

What $S(i;p)$ Are Proper?

Theorem. The scoring rule $S(i;p)$ is (strictly) proper if and only if there exists a (strictly) convex function $G: \Delta_n \rightarrow \mathbb{R}$ such that

$$S(i;p) = G(p) + \nabla G(p)(e_i - p)$$

Proof of "⇒"

- $S(\lambda;p) = \sum_{i\in[n]} \lambda_i S(i;p)$ is a linear fnc of λ for any p
- By properness, $S(\lambda;\lambda) = \max_{p\in\Delta_n} S(\lambda;p)$, denoted as $G(\lambda)$
  - $G(\lambda)$ is convex in λ since it is max of linear functions
- The gradient of $G(\lambda)$ is the gradient of $\sum_{i\in[n]} \lambda_i S(i;p)$ for the $p = \lambda$
  - I.e., $\nabla G(\lambda) = S(\cdot\,;\lambda)$
- Thus,

$$\begin{aligned} S(i;p) &= S(p;p) + [S(i;p) - S(p;p)] \\ &= G(p) + S(\cdot;p)\cdot[e_i - p] \\ &= G(p) + \nabla G(p)[e_i - p] \end{aligned}$$

What If There are Many Experts?

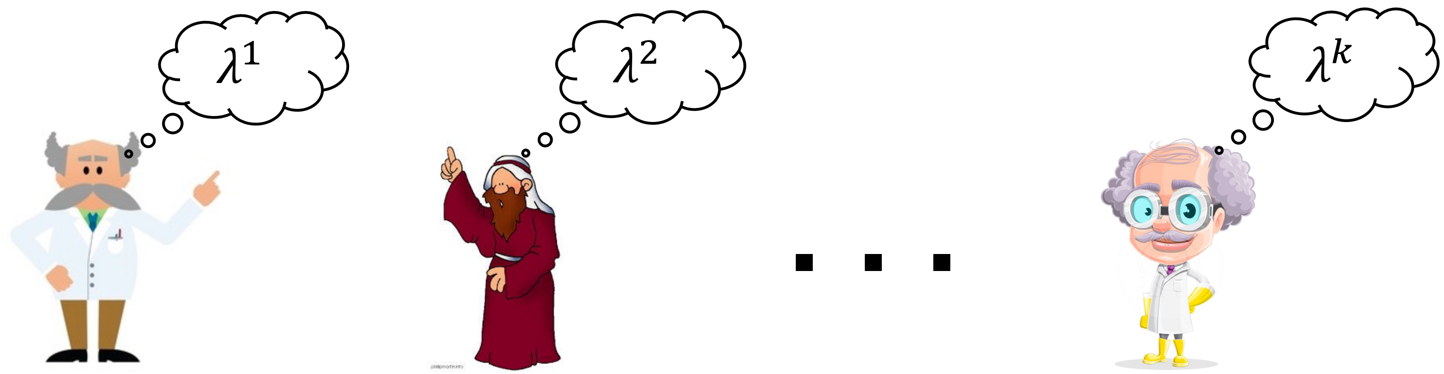

- One idea: elicit their predictions privately/separately
- Drawbacks
  1. May be expensive or wasteful – if experts all agree, we pay many times for the same prediction
  2. Not clear how to aggregate these predictions (average or geometric mean would not work)
  3. In fact, it may require experts' knowledge to correctly aggregate predictions

One Proper Way: Sequential Elicitation

- Ask experts to make predictions in sequence
- The reward for expert k's prediction p^k will be

$$S(i; p^k) - S(i; p^{k-1})$$

  where p^{k-1} is the prediction of expert $k-1$
  - I.e., experts are paid based on how much they improved the prediction

Theorem. If S is a proper scoring rule and each expert can only predict once, then each expert maximizes utility by reporting true belief given her own knowledge.

- Proof: since $S(i; p^{k-1})$ not under k's control, she maximizes reward by maximizing $S(i; p^k)$

One Proper Way: Sequential Elicitation

- Ask experts to make predictions in sequence
- The reward for expert k's prediction p^k will be

$$S(i; p^k) - S(i; p^{k-1})$$

where p^{k-1} is the prediction of expert $k-1$
  - I.e., experts are paid based on how much they improved the prediction

Theorem. If S is a proper scoring rule and each expert can only predict once, then each expert maximizes utility by reporting true belief given her own knowledge.

Remarks:

- k can see previous reports and then update his prediction
  - Experts will aggregate predictions automatically

One Proper Way: Sequential Elicitation

- Ask experts to make predictions in sequence
- The reward for expert k's prediction p^k will be

$$S(i; p^k) - S(i; p^{k-1})$$

where p^{k-1} is the prediction of expert $k-1$

  - I.e., experts are paid based on how much they improved the prediction

Theorem. If S is a proper scoring rule and each expert can only predict once, then each expert maximizes utility by reporting true belief given her own knowledge.

Remarks:

- Not true if an expert can report predictions for multiple times
  - She may manipulate her initial report to mislead others' prediction so that she has opportunity to significantly improve her prediction later
  - Will see an example later

One Proper Way: Sequential Elicitation

- Ask experts to make predictions in sequence
- The reward for expert k's prediction p^k will be

$$S(i; p^k) - S(i; p^{k-1})$$

where p^{k-1} is the prediction of expert $k-1$

  - I.e., experts are paid based on how much they improved the prediction

Theorem. If S is a proper scoring rule and each expert can only predict once, then each expert maximizes utility by reporting true belief given her own knowledge.

Q1: how does sequential elicitation relate to prediction market?

Q2: what happens is an expert can predict for multiple times?

Outline

- Scoring Rule and its Characterization
- Connection to Prediction Markets
- Gaming a Prediction Markets

Equivalence of PMs and Sequential Elicitation

Theorem (informal). Under mild technical assumptions, efficient prediction markets are in one-to-one correspondence to sequential information elicitation using proper scoring rules.

What does it mean?

- Experts will have exactly the same incentives and receive the same return
- Market maker's total loss = what elicitator's payment

Next: will *informally* illustrate using the LMSR and log-scoring rules

Equivalence of LMSR and Log-Scoring Rules

Recall LMSR

- Value function with current sales quantity q: $V(q) = b \ln \sum_{j \in [n]} e^{q_j/b}$
- To buy $x \in \mathbb{R}^n$ amount, a buyer pays: $V(q+x) - V(q)$
- Price function (they sum up to 1)

$$p_i(q) = \frac{e^{q_i/b}}{\sum_{j \in [n]} e^{q_j/b}} = \frac{\partial V(q)}{\partial q_i}$$

Fact. The optimal amount an expert purchases is the amount that moves the market price to her belief λ.

Fact. Worst case market maker loses is $b \ln n$.

Equivalence of LMSR and Log-Scoring Rules

Q1: If current market price is p^{k-1}, what is the optimal payoff for an expert with belief $\lambda = p^k$?

- Let q^{k-1} denote the market standing corresponding to price p^{k-1}
  - That is

$$\frac{e^{q_i^{k-1}/b}}{\sum_{j\in[n]} e^{q_j^{k-1}/b}} = p_i^{k-1}$$

Crucial terms:

- Value function $V(q) = b \ln \sum_{j\in[n]} e^{q_j/b}$
- Price function $p_i(q) = \frac{e^{q_i/b}}{\sum_{j\in[n]} e^{q_j/b}} = \frac{\partial V(q)}{\partial q_i}$

Equivalence of LMSR and Log-Scoring Rules

Q1: If current market price is p^{k-1}, what is the optimal payoff for an expert with belief $\lambda = p^k$?

- Let q^{k-1} denote the market standing corresponding to price p^{k-1}
- Optimal purchase for the expert is x^* *such that*

$$p_i\left(q^{k-1} + x^*\right) = \frac{e^{(q_i^{k-1}+x_i^*)/b}}{\sum_{j\in[n]} e^{(q_j^{k-1}+x_j^*)/b}} = p_i^k$$

and pays

$$V\left(q^{k-1} + x^*\right) - V(q^{k-1})$$

$$= b \ln \sum_{j\in[n]} e^{(q_j^{k-1}+x_j^*)/b} - b \ln \sum_{j\in[n]} e^{q_j^{k-1}/b}$$

Crucial terms:

- Value function $V(q) = b \ln \sum_{j\in[n]} e^{q_j/b}$
- Price function $p_i(q) = \frac{e^{q_i/b}}{\sum_{j\in[n]} e^{q_j/b}} = \frac{\partial V(q)}{\partial q_i}$

Equivalence of LMSR and Log-Scoring Rules

Q1: If current market price is p^{k-1}, what is the optimal payoff for an expert with belief $\lambda = p^k$?

- Let q^{k-1} denote the market standing corresponding to price p^{k-1}
- Optimal purchase for the expert is x^* *such that*

$$p_i(q^{k-1} + x^*) = \frac{e^{(q_i^{k-1}+x_i^*)/b}}{\sum_{j\in[n]} e^{(q_j^{k-1}+x_j^*)/b}} = p_i^k$$

and pays

$$V(q^{k-1} + x^*) - V(q^{k-1})$$

$$= b \ln \sum_{j\in[n]} e^{(q_j^{k-1}+x_j^*)/b} \;-\; b \ln \sum_{j\in[n]} e^{q_j^{k-1}/b}$$

$$\sum_{j\in[n]} e^{(q_j^{k-1}+x_j^*)/b} = \frac{e^{(q_i^{k-1}+x_i^*)/b}}{p_i^k}$$

Equivalence of LMSR and Log-Scoring Rules

Q1: If current market price is p^{k-1}, what is the optimal payoff for an expert with belief $\lambda = p^k$?

- Let q^{k-1} denote the market standing corresponding to price p^{k-1}
- Optimal purchase for the expert is x^* *such that*

$$p_i(q^{k-1} + x^*) = \frac{e^{(q_i^{k-1}+x_i^*)/b}}{\sum_{j\in[n]} e^{(q_j^{k-1}+x_j^*)/b}} = p_i^k$$

and pays

$$V(q^{k-1} + x^*) - V(q^{k-1})$$

$$= b \ln \sum_{j\in[n]} e^{(q_j^{k-1}+x_j^*)/b} - b \ln \sum_{j\in[n]} e^{q_j^{k-1}/b}$$

$$= b \ln \frac{e^{(q_i^{k-1}+x_i^*)/b}}{p_i^k} - b \ln \frac{e^{q_i^{k-1}/b}}{p_i^{k-1}}$$

Equivalence of LMSR and Log-Scoring Rules

Q1: If current market price is p^{k-1}, what is the optimal payoff for an expert with belief $\lambda = p^k$?

- Let q^{k-1} denote the market standing corresponding to price p^{k-1}
- Optimal purchase for the expert is x^* *such that*

$$p_i(q^{k-1} + x^*) = \frac{e^{(q_i^{k-1}+x_i^*)/b}}{\sum_{j\in[n]} e^{(q_j^{k-1}+x_j^*)/b}} = p_i^k$$

and pays

$$V(q^{k-1} + x^*) - V(q^{k-1})$$

$$= b\ln \sum_{j\in[n]} e^{(q_j^{k-1}+x_j^*)/b} - b\ln \sum_{j\in[n]} e^{q_j^{k-1}/b}$$

$$= b\ln \frac{e^{(q_i^{k-1}+x_i^*)/b}}{p_i^k} - b\ln \frac{e^{q_i^{k-1}/b}}{p_i^{k-1}}$$

$$= x_i^* - b(\ln p_i^k - \ln p_i^{k-1})$$

Note: this holds for any i

Equivalence of LMSR and Log-Scoring Rules

Q1: If current market price is p^{k-1}, what is the optimal payoff for an expert with belief $\lambda = p^k$?

- Let q^{k-1} denote the market standing corresponding to price p^{k-1}
- Record our finding: expert pays $x_i^* - b(\ln p_i^k - \ln p_i^{k-1})$
  - x^* is optimal amount for purchase
- What is the expert utility if outcome i is ultimately realized?

$$x_i^* - [x_i^* - b(\ln p_i^k - \ln p_i^{k-1})]$$

from contracts' return

Equivalence of LMSR and Log-Scoring Rules

Q1: If current market price is p^{k-1}, what is the optimal payoff for an expert with belief $\lambda = p^k$?

- Let q^{k-1} denote the market standing corresponding to price p^{k-1}
- Record our finding: expert pays $x_i^* - b(\ln p_i^k - \ln p_i^{k-1})$
  - x^* is optimal amount for purchase
- What is the expert utility if outcome i is ultimately realized?

$$x_i^* - [x_i^* - b(\ln p_i^k - \ln p_i^{k-1})]$$

$$= b \cdot [\ln p_i^k - \ln p_i^{k-1}]$$

$$= b \cdot [\, S^{log}(i; p^k) - S^{log}(i; p^{k-1})]$$

= reward in the sequential elicitation (up to scalar constant b)

Equivalence of LMSR and Log-Scoring Rules

Q1: If current market price is p^{k-1}, what is the optimal payoff for an expert with belief $\lambda = p^k$?

- Let q^{k-1} denote the market standing corresponding to price p^{k-1}
- Record our finding: expert pays $x_i^* \ - b(\ln p_i^k \ - \ln p_i^{k-1})$
  - x^* is optimal amount for purchase
- What is the expert utility if outcome i is ultimately realized?

Expert achieves the same utility in LMSR and log-scoring-rule elicitation *for any event realization*

Equivalence of LMSR and Log-Scoring Rules

Q2: What is the worst case loss (i.e., maximum possible payment) when using log-scoring rule in sequential info elicitation?

- Total payment – if event i realized – is

$$\sum_{k=1}^{K} [\ln p_i^k - \ln p_i^{k-1}] = \ln p_i^K - \ln p_i^0$$
$$\leq 0 - \ln p_i^0$$

- Start from $p^0 = (\frac{1}{n}, \cdots, \frac{1}{n})$ as the uniform distribution
- Worst-case loss is thus $\ln n$ (same as LMSR, up to constant b)

Back to Our Original Theorem…

Theorem (informal). Under mild technical assumptions, efficient prediction markets are in one-to-one correspondence with sequential information elicitation using proper scoring rules.

- Previous argument generalizes to arbitrary proper scoring rules
- Formal proof employs *duality theory*
  - Need something called "convex conjugate"

See paper *Efficient Market Making via Convex Optimization* for details

Outline

- Scoring Rule and its Characterization
- Connection to Prediction Markets
- Gaming a Prediction Markets

- Generally, we cannot force experts to participate just once
  - E.g., in prediction market, cannot force expert to just purchase once
- Manipulations arise when experts can predict multiple times
  - This is the case even two experts A, B and only A can predict twice
  - The so-called A-B-A game (arguably the most fundamental setting with multiple-round predictions)

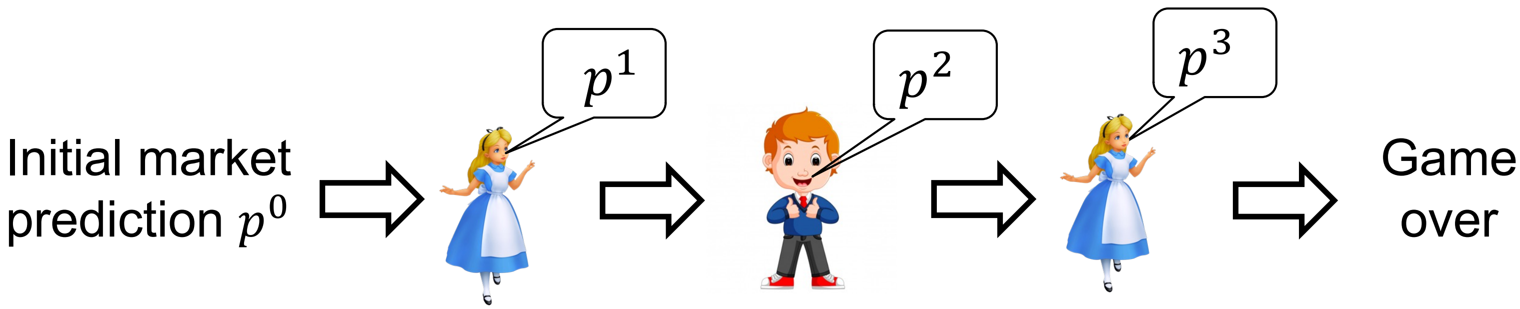

An Example of A-B-A Game

- Predict event $E \in \{0,1\}$; Outcome drawn uniformly at random
- Expert Alice observes a signal $A = E$
  - She exactly observes outcome
- Expert Bob also observes the outcome, i.e., signal $B = E$

Q: In A-B-A game, what should Alice predict at stage 1 and 3?

Report her true prediction at stage 1 (which is perfectly correct)

A-B-A Game: Example 2

- Alice observes signal $A \in \{0,1\}$, and $\Pr(A=0) = 0.51$
- Bob observes signal $B \in \{0,1\}$, and $\Pr(B=0) = 0.49$
  - A, B are independent
- They are asked to predict event $E = (\text{whether } A, B \text{ differ})$
  - The answer is YES or NO

Q: what is the optimal experts behaviors in A-B-A game?

Market starts with initial prediction $\mathrm{p}^0(\mathrm{YES}) = \mathrm{P}^0(\mathrm{NO}) = 1/2$

A-B-A Game: Example 2

- Alice observes signal $A \in \{0,1\}$, and $\Pr(A = 0) = 0.51$
- Bob observes signal $B \in \{0,1\}$, and $\Pr(B = 0) = 0.49$
  - A, B are independent
- They are asked to predict event $E = (\text{whether } A, B \text{ differ})$
  - The answer is YES or NO

Q: what is the optimal experts behaviors in A-B-A game?

- At stage 1, what is Alice's probability belief of YES?
  - If Alice's $A = 1$, then $\Pr(YES) = 0.49$
  - If Alice's $A = 0$, then $\Pr(YES) = 0.51$
- Should Alice report this at stage 1?
  - No, her truthful report tells B exactly the value of her A
  - Bob can then make a perfect prediction about E

A-B-A Game: Example 2

- Alice observes signal $A \in \{0,1\}$, and $\Pr(A=0) = 0.51$
- Bob observes signal $B \in \{0,1\}$, and $\Pr(B=0) = 0.49$
  - A, B are independent
- They are asked to predict event $E = (\text{whether } A, B \text{ differ})$
  - The answer is YES or NO

Q: what is the optimal experts behaviors in A-B-A game?

- What should Alice do at stage 1 then?
  - Say nothing, or equivalently, predict $p^1 = p^0$

A-B-A Game: Example 2

- Alice observes signal $A \in \{0,1\}$, and $\Pr(A = 0) = 0.51$
- Bob observes signal $B \in \{0,1\}$, and $\Pr(B = 0) = 0.49$
  - A, B are independent
- They are asked to predict event $E = (\text{whether } A, B \text{ differ})$
  - The answer is YES or NO

Q: what is the optimal experts behaviors in A-B-A game?

- What should Bob predict at stage 2?
  - Bob learns nothing from stage 1
  - So If $B = 1$, then $\Pr(YES) = 0.51$; if $B = 0$, then $\Pr(YES) = 0.49$
  - Should report truthfully based on the above belief – why?

He only has one chance to predict, and his belief is the best given his current knowledge

A-B-A Game: Example 2

- Alice observes signal $A \in \{0,1\}$, and $\Pr(A = 0) = 0.51$
- Bob observes signal $B \in \{0,1\}$, and $\Pr(B = 0) = 0.49$
  - A, B are independent
- They are asked to predict event $E = (\text{whether } A, B \text{ differ})$
  - The answer is YES or NO

Q: what is the optimal experts behaviors in A-B-A game?

- What should Bob predict at stage 2?
  - Bob learns nothing from stage 1
  - So If $B = 1$, then $\Pr(YES) = 0.51$; if $B = 0$, then $\Pr(YES) = 0.49$
  - Should report truthfully based on the above belief – why?
  - Bob's truthful report reveals his signal, but gains little utility

A-B-A Game: Example 2

- Alice observes signal $A \in \{0,1\}$, and $\Pr(A = 0) = 0.51$
- Bob observes signal $B \in \{0,1\}$, and $\Pr(B = 0) = 0.49$
  - A, B are independent
- They are asked to predict event $E = (\text{whether } A, B \text{ differ})$
  - The answer is YES or NO

Q: what is the optimal experts behaviors in A-B-A game?

- What should Alice predict at stage 3?
  - She just learned Bob's signal B
  - So can precisely predict "whether A, B differ" now
  - Alice now moves the prediction from $\Pr(YES) = 0.51 \; or \; 0.49$ to $\Pr(YES) = 1 \; or \; 0$ → receiving a lot of credits

A-B-A Game: Example 2

- Alice observes signal $A \in \{0,1\}$, and $\Pr(A = 0) = 0.51$
- Bob observes signal $B \in \{0,1\}$, and $\Pr(B = 0) = 0.49$
  - A, B are independent
- They are asked to predict event $E = (\text{whether } A, B \text{ differ})$
  - The answer is YES or NO

Remarks

- Example shows how experts aggregate previous information and update their predictions along the way
- Gaming behaviors arise even if a single expert can predict twice

A-B-A Game: Example 2

- Alice observes signal $A \in \{0,1\}$, and $\Pr(A = 0) = 0.51$
- Bob observes signal $B \in \{0,1\}$, and $\Pr(B = 0) = 0.49$
  - A, B are independent
- They are asked to predict event $E = (\text{whether } A, B \text{ differ})$
  - The answer is YES or NO

Remarks

- This is an issue in prediction markets, since experts can buy and sell whenever they want
- Equilibrium of PMs are still poorly understood, even for the fundamental A-B-A games
  - See a recent paper *Computing Equilibria of Prediction Markets via Persuasion*

Thank You

Haifeng Xu

University of Chicago

haifengxu@uchicago.edu