CS6501: Topics in Learning and Game Theory (Fall 2019)

Prediction Markets and Scoring Rules

Instructor: Haifeng Xu

Outline

- Recap: Scoring Rule and Information Elicitation
- Connection to Prediction Markets
- Manipulations in Prediction Markets

Information Elicitation from A Single Expert

- We (designer) want to learn the distribution of random var $E \in [n]$
  - E will be sampled in the future
- An expert/predictor has a predicted distribution $\lambda \in \Delta_n$
- Want to incentivize the expert to truthfully report λ

Idea: reward expert by designing a scoring rule $S(i; p)$ where:
(1) p is the expert's report (may not equal λ);
(2) $i \in [n]$ is the event realization

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

Proper Scoring Rules

Example 1 [Log Scoring Rule]

- $S(i; p) = \log p_i$

Example 2 [Quadratic Scoring Rule]

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

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

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

basis vector

Information Elicitation from Many Experts

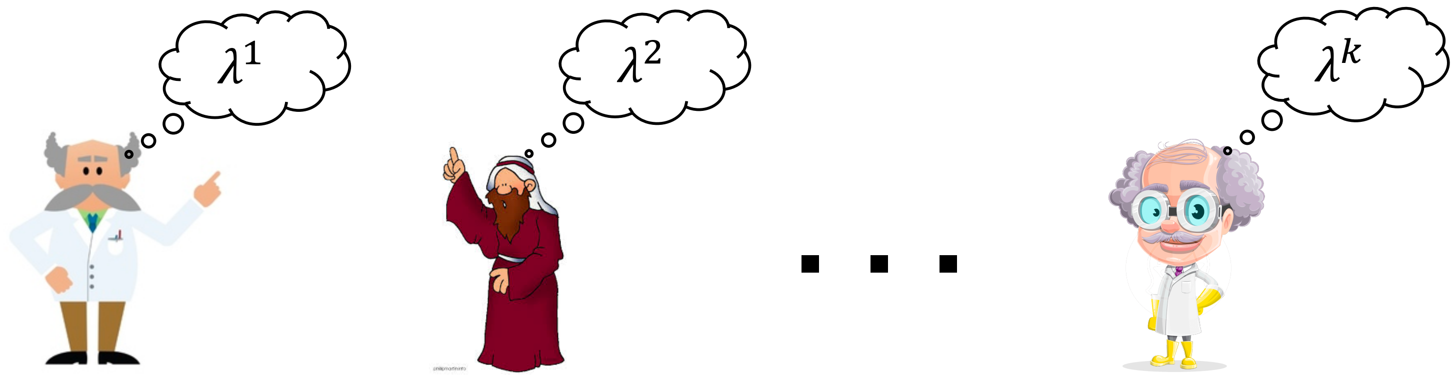

Idea: sequential elicitation – experts make predictions in sequence

- Reward for expert k's prediction p^k is

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

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

Information Elicitation from Many Experts

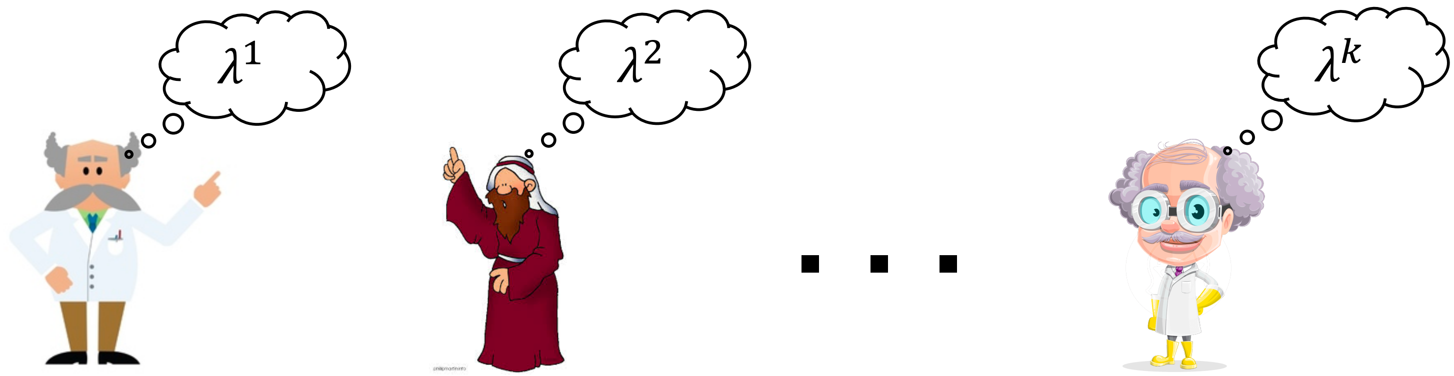

Theorem. If S is a proper scoring rule and each expert can only predict once, then each expert maximizes utility by reporting true belief given her own knowledge.

Remark

- ➢Each expert is expected to improve the prediction by aggregating previous predictions and then update it
  - Otherwise they will lose money

Information Elicitation from Many Experts

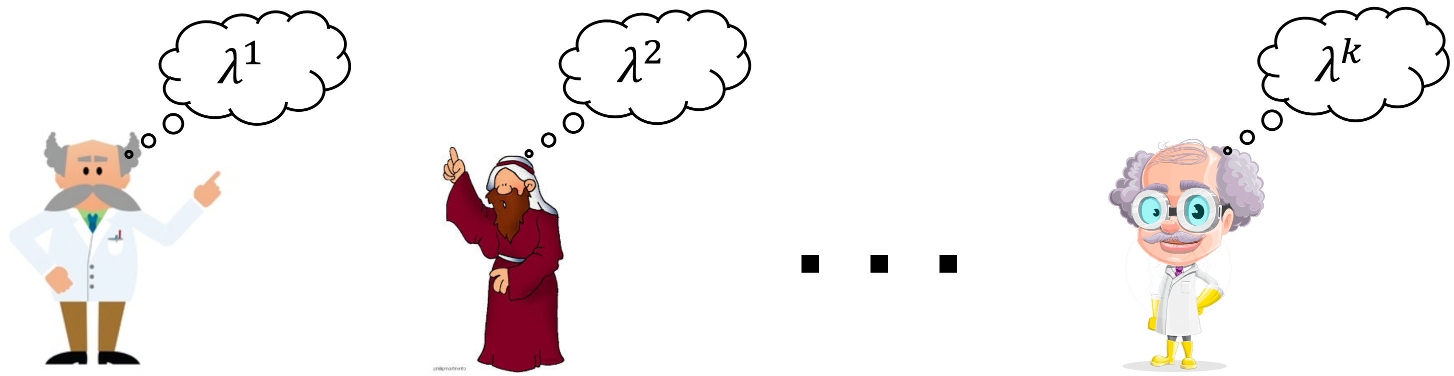

Theorem. If S is a proper scoring rule and each expert can only predict once, then each expert maximizes utility by reporting true belief given her own knowledge.

Q1: how does sequential elicitation relate to prediction market?

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

Outline

- Recap: Scoring Rule and Information Elicitation
- Connection to Prediction Markets
- Manipulations in 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 is the same

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

Equivalence of LMSR and Log-Scoring Rules

Recall LMSR

- Value function with current sales quantity q: $V(q) = b \log \sum_{j\in[n]} e^{q_j/b}$
- 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 $b \log n$.

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 $V(q) = b \log \sum_{j\in[n]} e^{q_j/b}$
- 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

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

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

Crucial terms:

- Value function $V(q) = b \log \sum_{j\in[n]} e^{q_j/b}$
- 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*

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

and pays

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

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

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

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*

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

and pays

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

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

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

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*

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

and pays

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

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

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

Note: this holds for any i

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

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}
- Repeat our finding: expert pays $x_i^* - b(\log p_i^k - \log p_i^{k-1})$
  - x^* is optimal amount for purchase
- What is the expert utility if outcome i is ultimately realized?

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

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}
- Repeat our finding: expert pays $x_i^* - b(\log p_i^k - \log p_i^{k-1})$
  - x^* is optimal amount for purchase
- What is the expert utility if outcome i is ultimately realized?

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

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

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

= payment in the sequential elicitation (constant b is a scalar)

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}
- Repeat our finding: expert pays $x_i^* - b(\log p_i^k - \log p_i^{k-1})$
  - 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

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

- To avoid cases where some p_i^0 is too small (then we need to pay a lot), should choose $p^0 = (\frac{1}{n}, \cdots, \frac{1}{n})$ as uniform distribution
- Worst-case loss is thus $\log n$ (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*
  - Recall, any proper scoring rule corresponds to a convex function
  - A prediction market is determined by a value function $V(q)$

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*
  - Recall, any proper scoring rule corresponds to a convex function
  - A prediction market is determined by a value function $V(q)$

The Correspondence

PM with $V(q)$ corresponds to sequential elicitation with scoring rules determined by $V^*(p) =$ the convex conjugate of $V(q)$

- Convex conjugate is in some sense the "dual" of function $V(q)$
- See paper *Efficient Market Making via Convex Optimization* for details

Outline

- Recap: Scoring Rule and Information Elicitation
- Connection to Prediction Markets
- Manipulations in 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)

Initial market prediction p^0

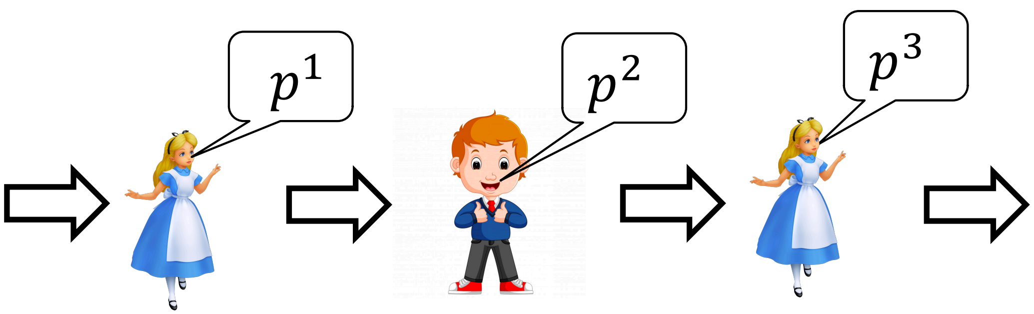

Game over

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 $E = (\text{whether } A + B = 1)$
  - 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 $E = (\text{whether } A + B = 1)$
  - 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

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 $E = (\text{whether } A + B = 1)$
  - 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 $E = (\text{whether } A + B = 1)$
  - 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 $E = (\text{whether } A + B = 1)$
  - 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 $E = (\text{whether } A + B = 1)$
  - 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 $A + B = 1$" 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 $E = (\text{whether } A + B = 1)$
  - The answer is YES or NO

Remarks

- Example shows how experts aggregate previous information and update their predictions along the way
- Manipulations 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 $E = (\text{whether } A + B = 1)$
  - 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* for state-of-the-art results

Thank You

Haifeng Xu

University of Virginia

hx4ad@virginia.edu