Two Strongly Truthful Mechanisms for Three Heterogeneous Agents Answering One Question

Grant Schoenebeck, **Fang-Yi Yu**

MICHIGAN UNIVERSITY OF MICHIGAN

Harvard John A. Paulson School of Engineering and Applied Sciences

Information elicitation without verification

- Subjective
  - Are you happy?
  - Do you like the restaurant?

Elicit Information from Crowds

- Agents' signals are dependent.
- Peer's report can elicit agents' truthful reports.
- Agents can report
  - Own signal
  - Prediction on other's signal
  - …

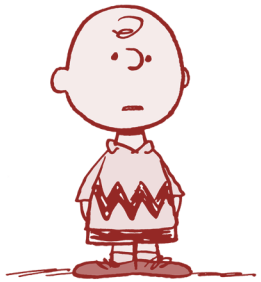

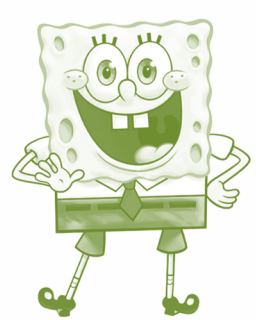

Contribution

We propose two differential peer prediction mechanisms

- Strongly truthful
  - Truth-telling is strict Bayesian Nash Equilibrium
  - and pays strictly higher than any other equilibria*
- The guarantees hold for
  - three agents
  - single item report
  - asymmetric priors
  - prior free

Outline

- Proper scoring rule with verification
  - Log scoring rule and mutual information
- Differential Peer Prediction Mechanisms
  - Proper scoring rule in IEWV
  - Source-DPP
  - Target-DPP
- Connection to BTS

Prediction with Verification

- (W, A) is sampled from a joint distribution P
  - $W \in \Omega$, ground truth
  - A, Alice's signal
- Proper scoring rule $S(w, p) \in \mathbb{R}$
  - $W = w$
  - $p \in \Delta^{\Omega}$, prediction
  - Alice maximizes $S(p, w)$ by honestly reporting $P(W \mid A = a)$

S is committed

Alice observes $A = a$

Alice reports prediction p

$W = w$ is revealed.

Alice gets $S(w, p)$

Log scoring rule $LS(p, w) = \log p(w)$

Three properties

- Proper: truth-telling $P_{W|A}$ maximizes $LS(p, w)$
- Shannon mutual information: $\mathbb{E}_{W,A}\left[LS\left(P_{W|A}, W\right)\right] = MI(W; A)$

Multiple predictions

- $W, A_1, A_2, \ldots$ is sampled from a joint distribution P
- Truth-telling $p_i = P_{W|(A_1,\ldots A_i)}$
- Chain rule
  $\mathbb{E}[LS(p_i, W) - LS(p_{i-1}, W)]$
  $= MI(W; A_i \mid A_1, \ldots, A_{i-1})$

S is committed

Alice observes $A_1 = a_1$ and reports p_1

…

Alice observes $A_i = a_i$ and reports p_i

Properties of Log scoring rule

- Proper
  - truth-telling $P_{W|A}$ maximizes $LS(p, w)$
- Chain rule [KS19]
  - $\mathbb{E}[LS(p_i, W) - LS(p_{i-1}, W)] = MI(W; A_i \mid A_1, \ldots, A_{i-1})$

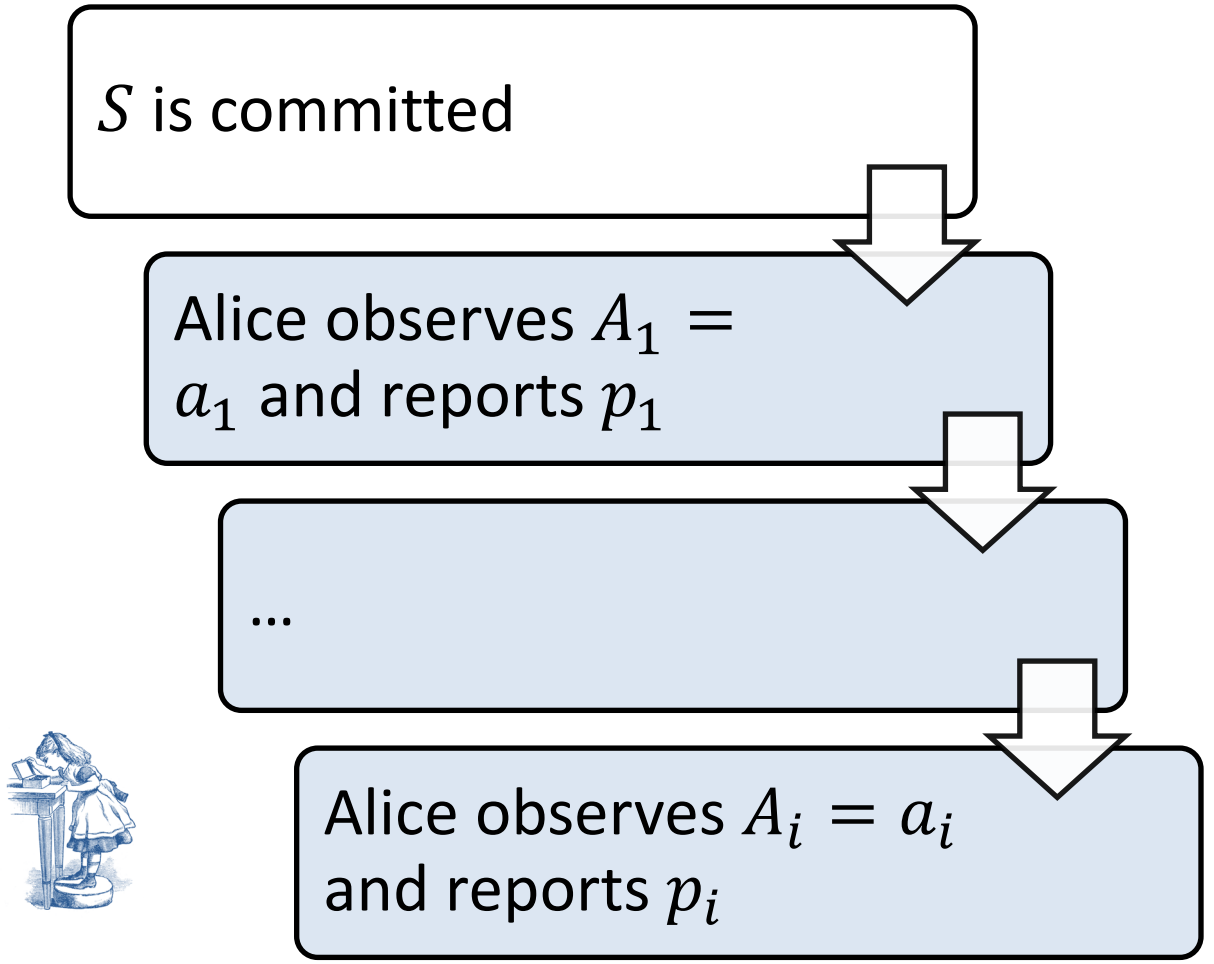

Outline

- Proper scoring rule with verification
  - Log scoring rule and mutual information
- Differential Peer Prediction Mechanisms
  - Proper scoring rule in IEWV
  - Source-DPP
  - Target-DPP
- Connection to BTS

Proper Scoring Rule $S(w, p)$ in IEWV

- Agents can play one of three roles
  - Target reports his signal and is predicted.
  - Expert makes prediction the target's signal.
  - Source provides information to the expert.

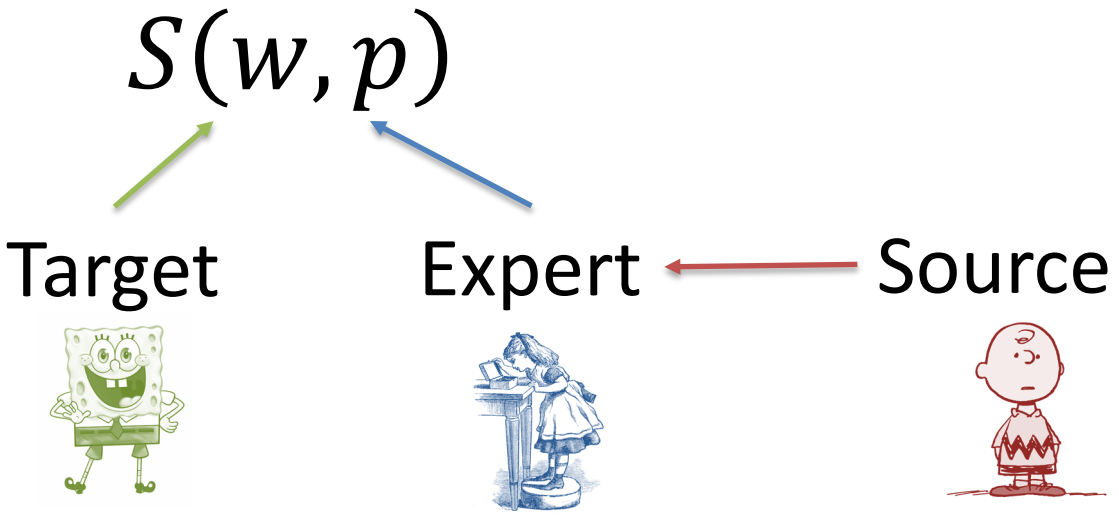

Outline

- Proper scoring rule with verification
  - Log scoring rule and mutual information
- Differential Peer Prediction Mechanisms
  - Proper scoring rule in IEWV
  - Source-DPP
  - Target-DPP
- Connection to BTS

Differential Peer Prediction Mechanisms

Bob and Charlie report their signals

Set one as Target and the other as the Source randomly.
Set Alice as the expert.

Alice reports her initial prediction Q on **Target's** signal.

Alice learns **Source's** signal.

Alice updates her improved prediction Q^+ on **Target's** signal.

Outline

- Proper scoring rule with verification
  - Log scoring rule and mutual information
- Differential Peer Prediction Mechanisms
  - Proper scoring rule in IEWV
  - Source-DPP
  - Target-DPP
- Connection to BTS

Source Differential Peer Prediction

Procedure

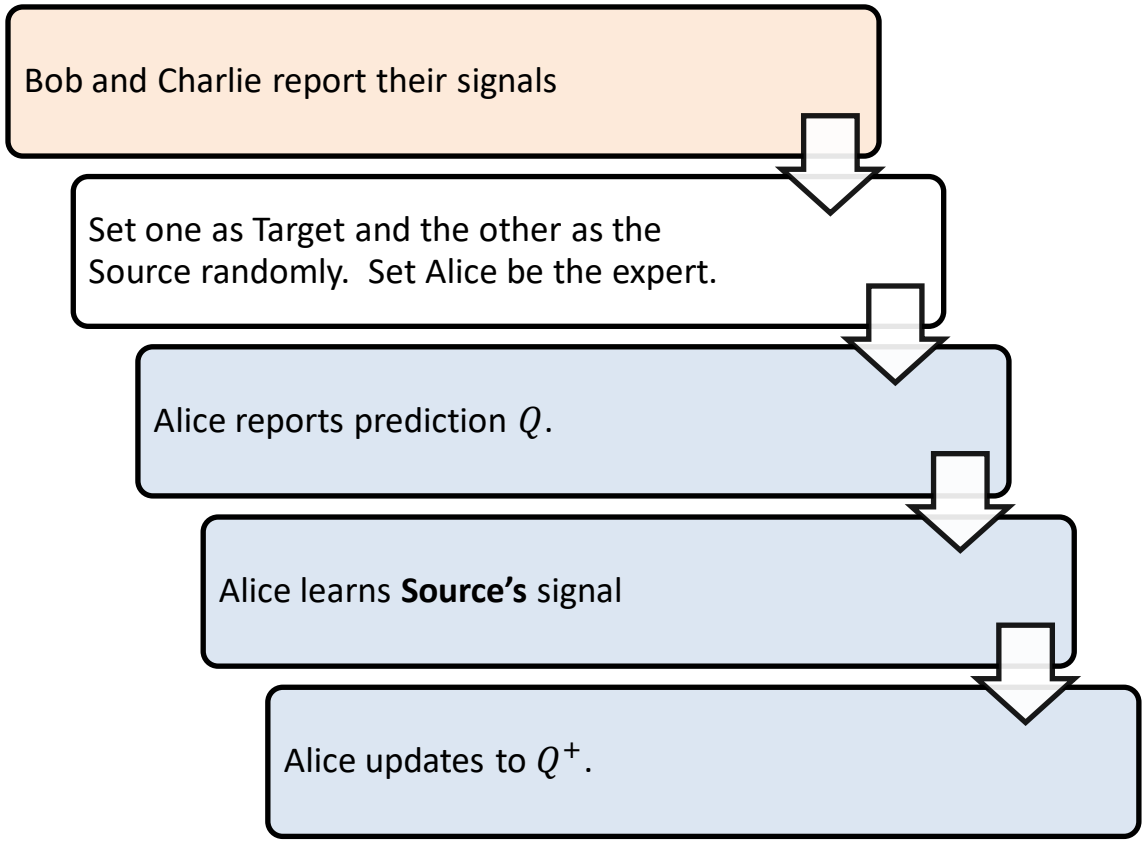

Payment (if Bob is the Target)

- Bob gets zero.
- Alice gets sum of scores,
$$LS(Q^+, B) + LS(Q, B)$$
- Charlie gets the improved score,
$$LS(Q^+, B)$$

Source wants to improve the expert's prediction

Truthful report from Expert and Source

- Agents can play one of three roles
  - Target reports his signal and is predicted.
  - Expert makes prediction the target's signal.
  - Source provides information to the expert.

log scoring rule is proper

$$S(w, p)$$

Target

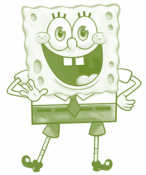

Expert

Source

Strongly truthfulness

Procedure

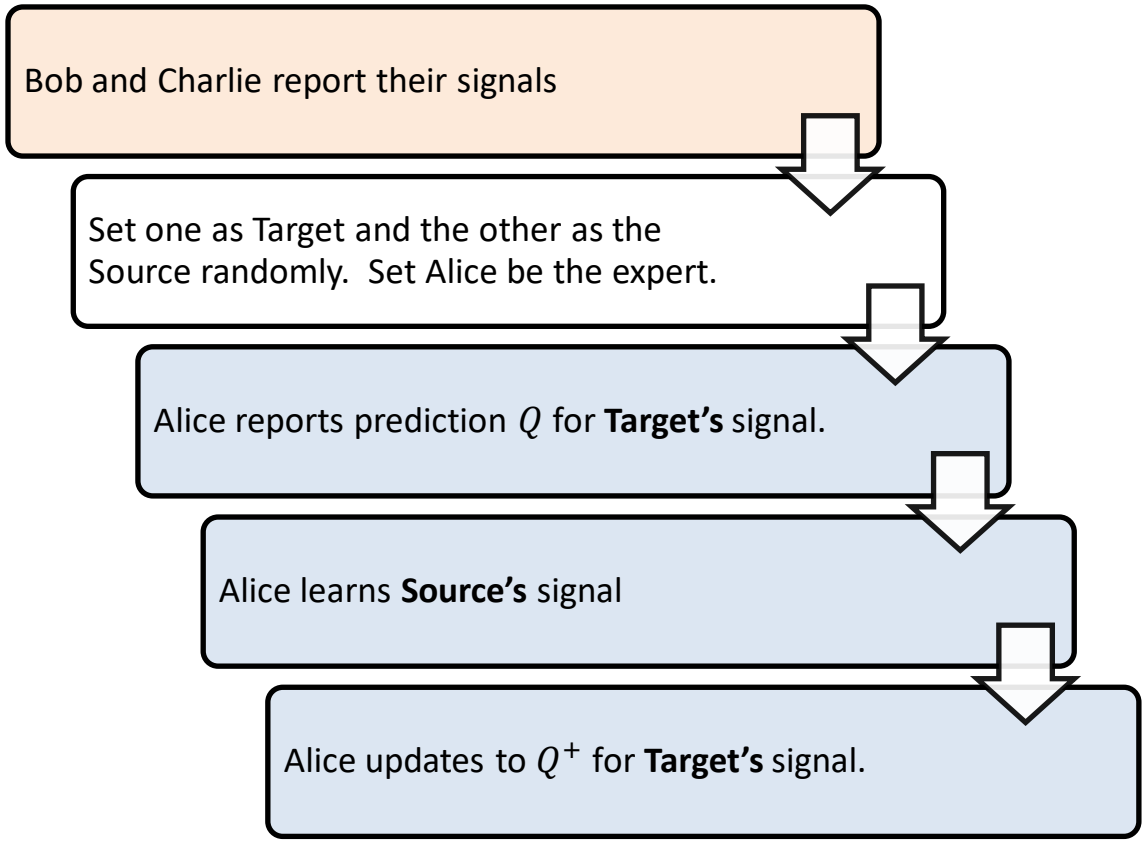

Payment (if Bob is the Target)

- Bob gets zero.
- Alice gets sum of scores,
  $$LS(Q^+, B) + LS(Q, B)$$
- Charlie gets the difference,
  $$LS(Q^+, B) - 3LS(Q, B)$$
- Total payment
  $$2(LS(Q^+, B) - LS(Q, B))$$
  whose expectation is $2MI(C; B \mid A)$

Outline

- Proper scoring rule with verification
  - Log scoring rule and mutual information
- Differential Peer Prediction Mechanisms
  - Proper scoring rule in IEWV
  - Source-DPP
  - Target-DPP
- Connection to BTS

Truthful report from Target

- Agents can play one of three roles
  - Target reports his signal and is predicted.
  - Expert makes prediction the target's signal.
  - Source provides information to the expert.

log scoring rule is proper

$$LS(w, p)$$

Target

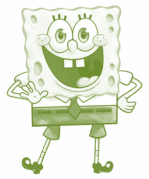

Expert

Source

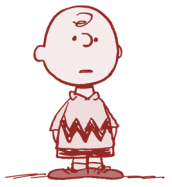

Truthful report from Target

- Agents can play one of three roles
  - Target reports his signal and is predicted.
  - Expert makes prediction the target's signal.
  - Source provides information to the expert.
- Proper scoring rule
  - Truthfulness the second argument

$$LS(w, p)$$

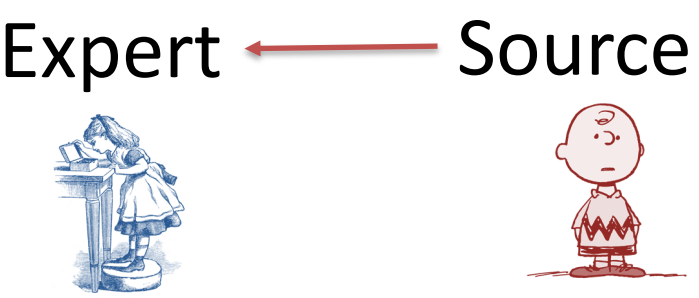

Truthful report from Target

- Agents can play one of three roles
  - Target reports his signal and is predicted.
  - Expert makes prediction the target's signal.
  - Source provides information to the expert.
- Proper scoring rule
  - Truthfulness the second argument
  - First argument?

$$LS(w, p)$$

Target

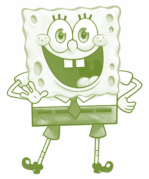

Reversed Log scoring rule

- Symmetry between Target and Source

$$\mathbb{E}[LS(Q^+, B) - LS(Q, B)] = MI(C; B \mid A)$$

- Reversed Log scoring rule

$$R(w) = LS(Q^+, w) - LS(Q, w)$$

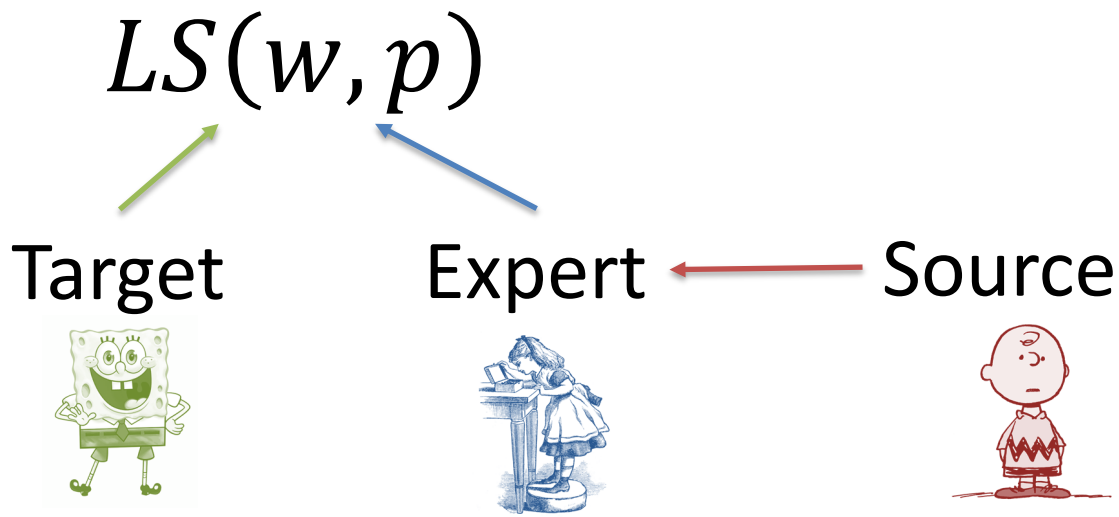

Target Differential Peer Prediction

Procedure

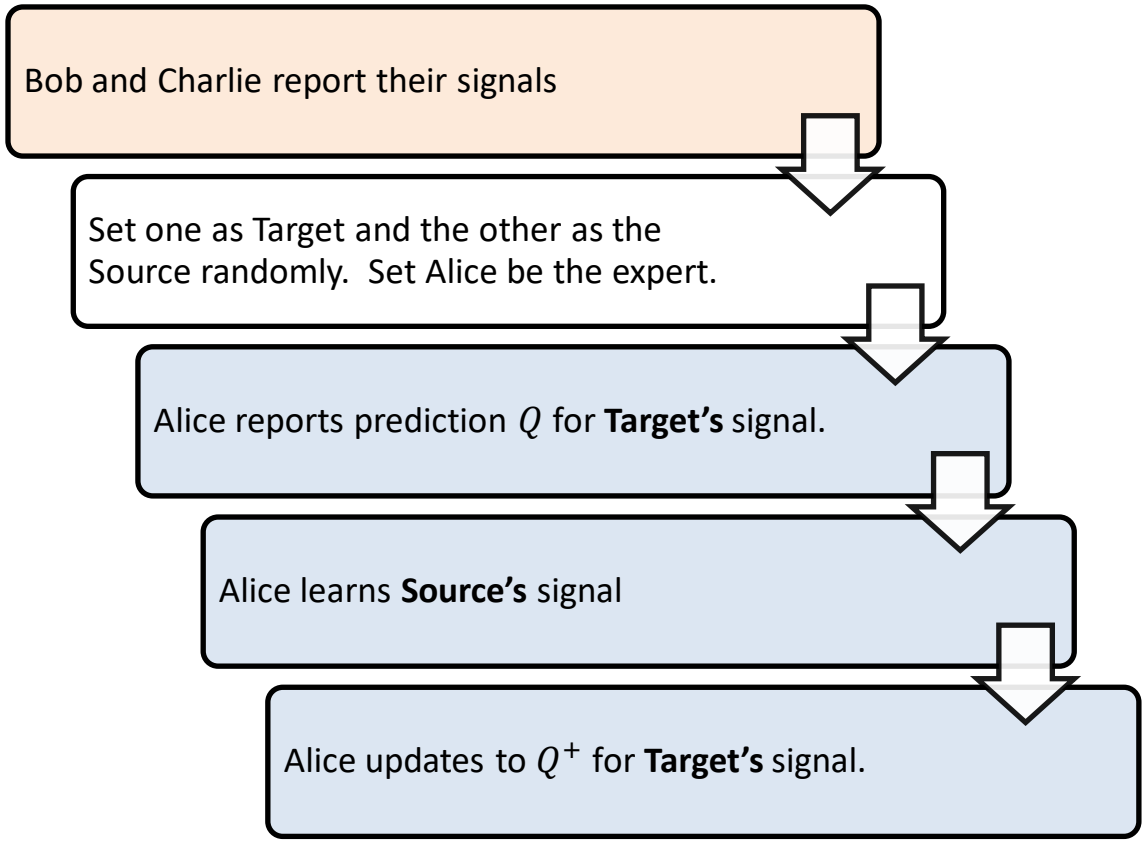

Payment (if Bob is the Target)

- Alice gets sum of scores,
  $$LS(Q^+, B) + LS(Q, B)$$
- Bob gets the difference,
  $$LS(Q^+, B) - LS(Q, B)$$
  whose expectation is $MI(C; B \mid A)$

Target wants to get predicted well.

Strongly truthfulness

Procedure

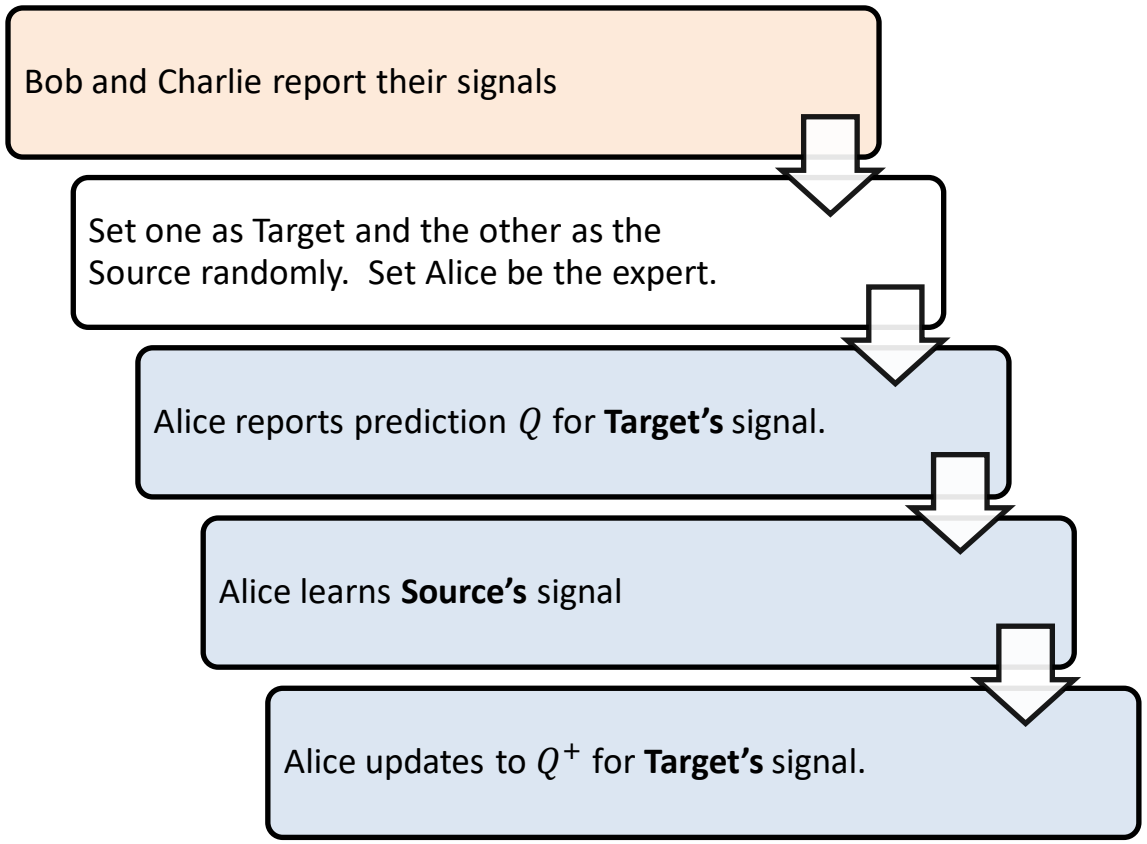

Payment (if Bob is the Target)

- Alice gets sum of scores,
  $$LS(Q^+, B) + LS(Q, B)$$
- Bob gets the difference,
  $$LS(Q^+, B) - LS(Q, B)$$
- Charlie gets
  $$-2LS(Q, B)$$
- Total payment
  $$2(LS(Q^+, B) - LS(Q, B))$$
  whose expectation is $2MI(C; B \mid A)$

Comparison

Target incentive

- Use Target's (Bob's) report as the ground truth
- Pro and Con
  - Log scoring rule only*
  - Parallel process

Source incentive

- Use Source's (Charlie's) report to improve prediction $Q \rightarrow Q^+$
- Pro and Con
  - works for all scoring rule
  - sequential process

Payment	Target DPP	Source DPP
Alice (Expert)	$LS(Q^+, B) + LS(Q, B)$	$LS(Q^+, B) + LS(Q, B)$
Bob (Target)	$LS(Q^+, B) - LS(Q, B)$	0
Charlie (Source)	$-2LS(Q, B)$	$LS(Q^+, B) - 3LS(Q, B)$
Total	$2(LS(Q^+, B) - LS(Q, B))$	$2(LS(Q^+, B) - LS(Q, B))$

Outline

- Proper scoring rule with verification
  - Log scoring rule and mutual information
- Differential Peer Prediction Mechanisms
  - Proper scoring rule in IEWV
  - Source-DPP
  - Target-DPP
- Connection to BTS

Related Work in Peer Prediction

Target-incentive mechanism

- Bayesian Truth Serum
  - Each i agent reports her signal x_i and prediction p_i on other's signal
  - Prediction score: measure the quality the prediction p_i
  - Information score: $LS(Q^+, x_i) - LS(Q, x_i)$

Aggregated prediction One agent's prediction

Source-incentive mechanism

- Robust BTS [Witkowski,Parkes 2011]
- Knowledge-Free Peer Prediction [Zhang, Chen 2014]

Question and Discussion