Optimal Scoring Rule Design under Partial Knowledge

Yiling Chen, Harvard University

Fang-Yi Yu, George Mason University

Harvard John A. Paulson School of Engineering and Applied Sciences

GEORGE MASON UNIVERSITY

High quality information form crowd

- Peer review at conferences
- Peer grading in classrooms
- Expert forecasting and predictions

- ○ **1.** Strong Reject 5%
- ○ **2.** Round 1 Reject 50%
- ○ **3.** Probable Eventual Reject 65%
- ○ **4.** Borderline (avoid using if possible) 70%
- ◉ **5.** Weak Accept 80 %
- ○ **6.** Accept 90%
- ○ **7.** Strong Accept 95%
- ○ **8.** Top (Best Paper Nomination) 99%
- ○ **9.** Very Top (Best Paper) 100%

Outline

- Problem set up
  - Proper scoring rule
  - Maximizing information gain
  - Partial knowledge
- Savage characterization
- Main results
  - Core idea: Information gain = convexity
  - Simulation

Incentivize predictions

- Peer review example
  - Binary outcome: $w \in \{0,1\}$
  - An agent's review/prediction: $x \in [0,1]$

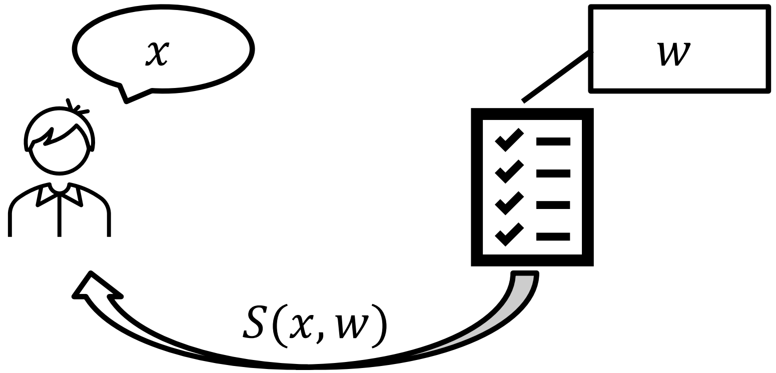

Incentivize predictions

- Peer review example
  - Binary outcome: $w \in \{0,1\}$
  - An agent's review/prediction: $x \in [0,1]$
- A scoring rule S rewards $S(x,w) \in \mathbb{R}$
  - S is proper if for all x'
    $\mathbb{E}_{w\sim x}[S(\boldsymbol{x},w)] \geq \mathbb{E}_{w\sim x}[S(\boldsymbol{x'},w)]$.

    Truthful Non-truthful

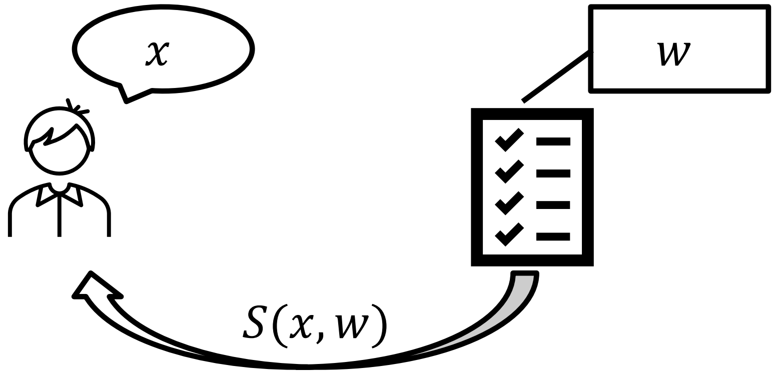

Incentivize predictions

- Peer review example
  - Binary outcome: $w \in \{0,1\}$
  - An agent's review/prediction: $x \in [0,1]$
- A scoring rule S rewards $S(x, w) \in \mathbb{R}$
  - S is proper if for all x'

    $\mathbb{E}_{w\sim x}[S(\boldsymbol{x}, w)] \geq \mathbb{E}_{w\sim x}[S(\boldsymbol{x}', w)].$

    $S(x, x) \geq S(x', x).$

    Truthful Non-truthful

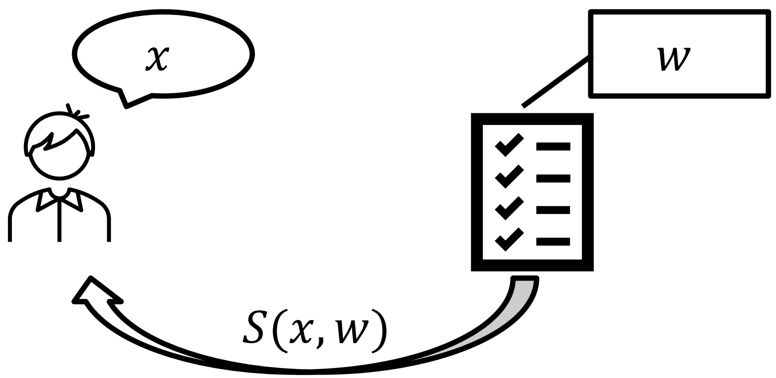

Is proper scoring rule enough?

- Peer review
  1. Principle announce S
  2. Agent reports $x \in [0,1]$
  3. Outcome reveals $w \in \{0,1\}$
  4. Agent gets $S(x, w)$

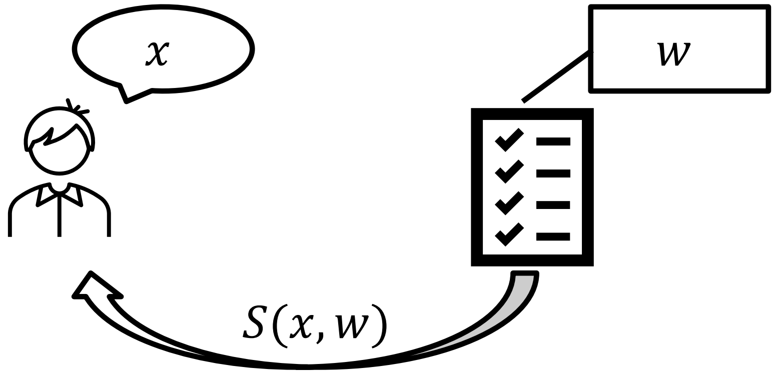

Incentivize costly predictions

- Peer review with effort
  1. Principle announce S
  2. Agent decides to acquire costly information
  3. Agent reports $x \in [0,1]$
  4. Outcome reveals $w \in \{0,1\}$
  5. Agent gets $S(x, w)$

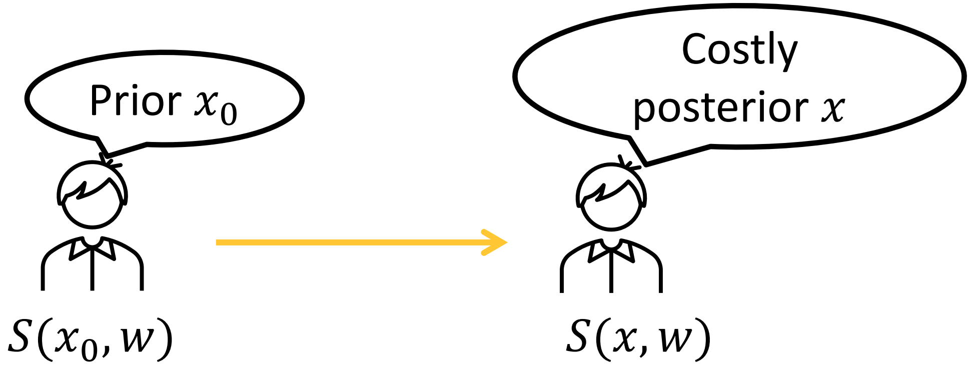

Incentivize costly predictions

- Peer review with effort
  1. Principle announce S
  2. Agent decides to acquire costly information
  3. Agent reports $x \in [0,1]$
  4. Outcome reveals $w \in \{0,1\}$
  5. Agent gets $S(x, w)$

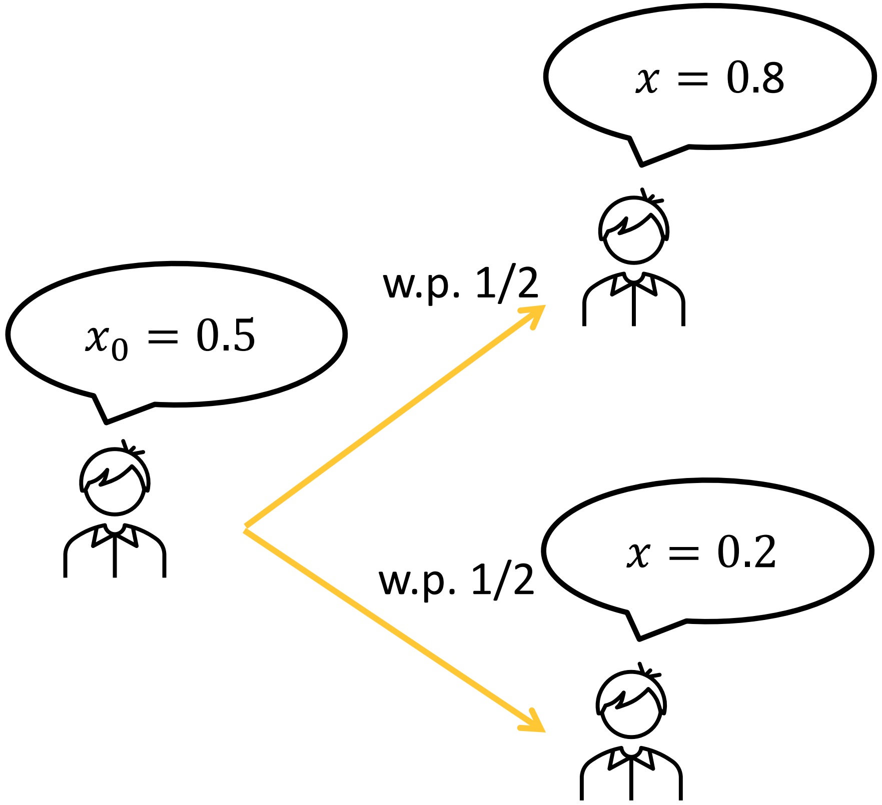

Optimization of scoring rule [HLSW20]

- Given an information structure P on (w, x), design bounded S so that maximize the expected gain

$$\max_S \text{Expected gain}$$
$$\text{s.t. } S \text{ is proper and bounded}$$

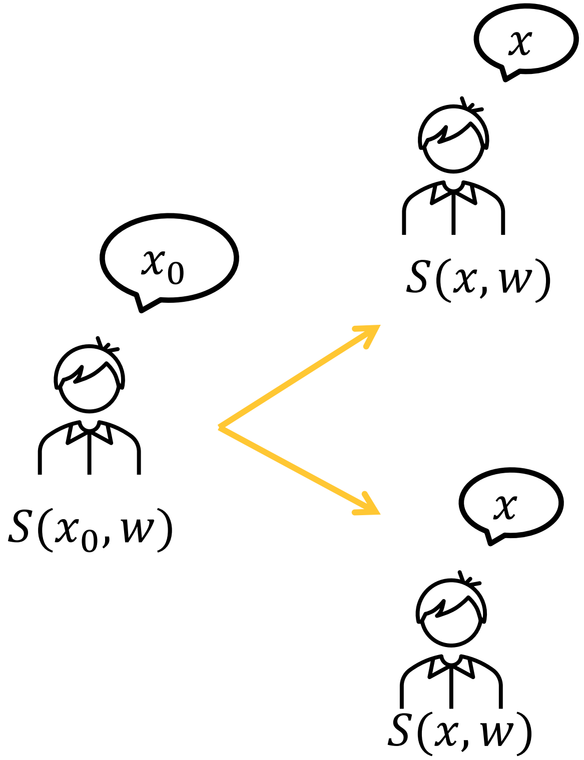

Optimization of scoring rule

- Given an information structure P on (w, x), design bounded S so that maximize the expected gain

$$\max_S \mathbb{E}_P[S(x,w) - S(x_0,w)]$$
$$\text{s.t. } S \text{ is proper and bounded}$$

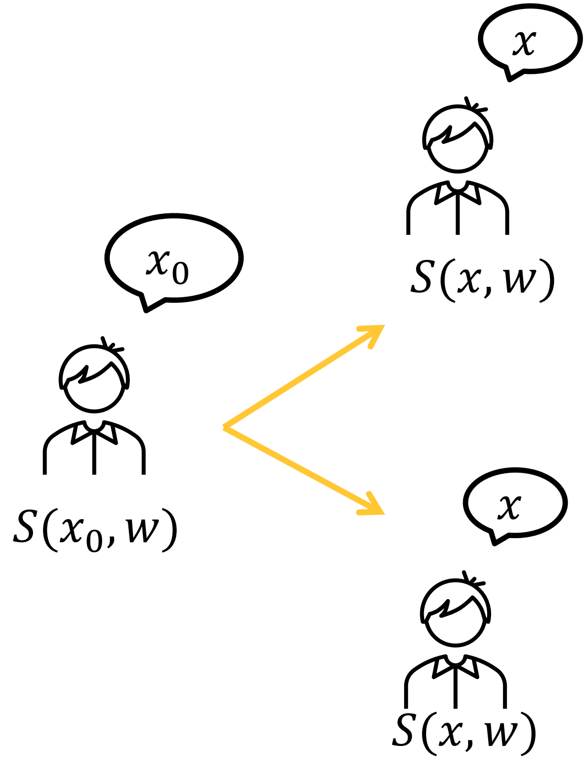

Multiple possible information structure P

Heterogeneous agents $\mathcal{P} = \{P_1, P_2 \ldots\}$

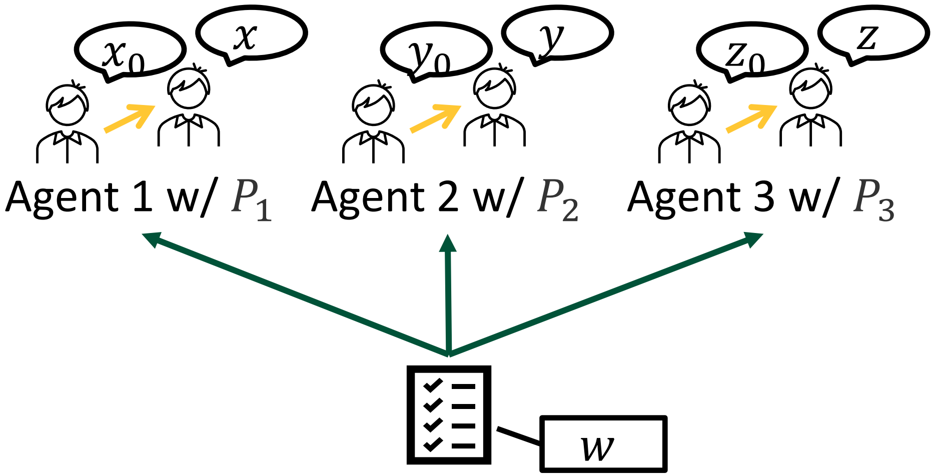

Multiple possible information structure P

Heterogeneous agents $\mathcal{P} = \{P_1, P_2 \ldots\}$

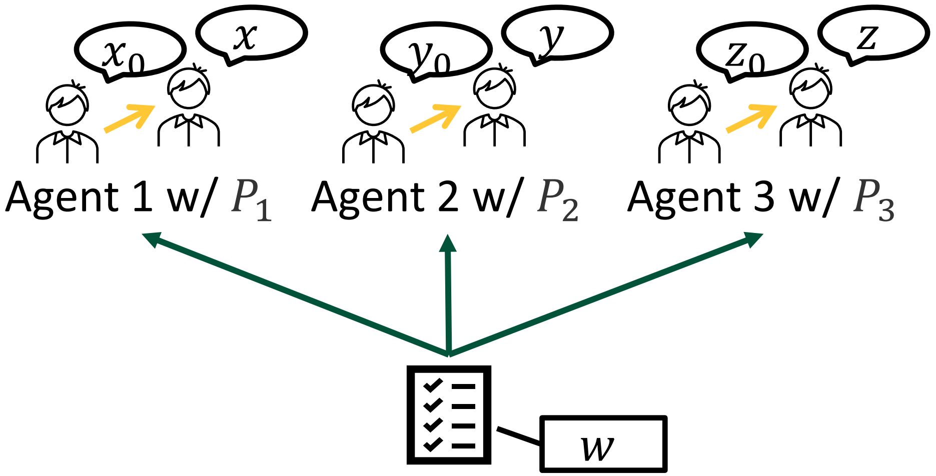

Sequential learning

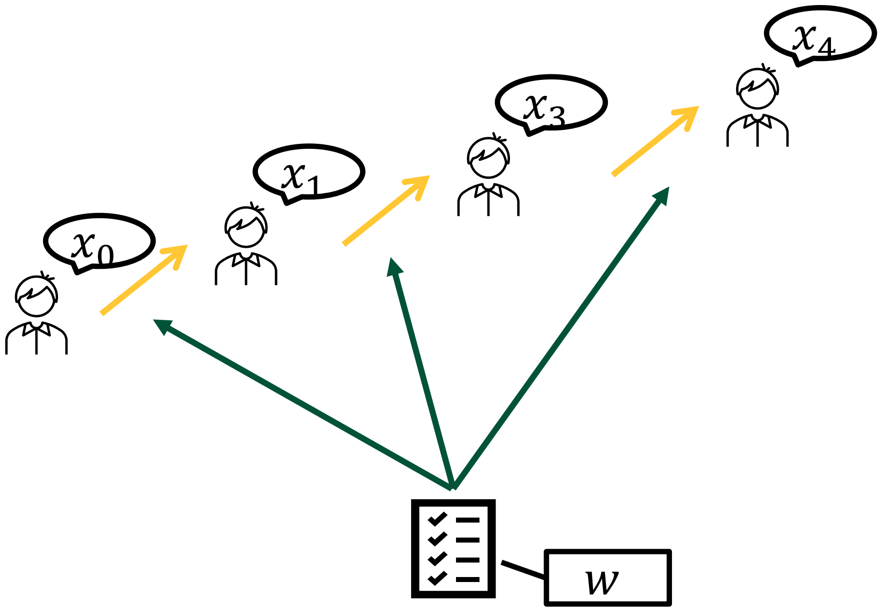

Multiple possible information structure P

Heterogeneous agents $\mathcal{P} = \{P_1, P_2 \ldots\}$

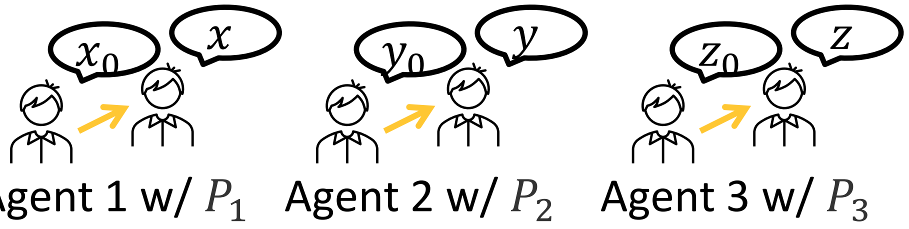

Sequential learning

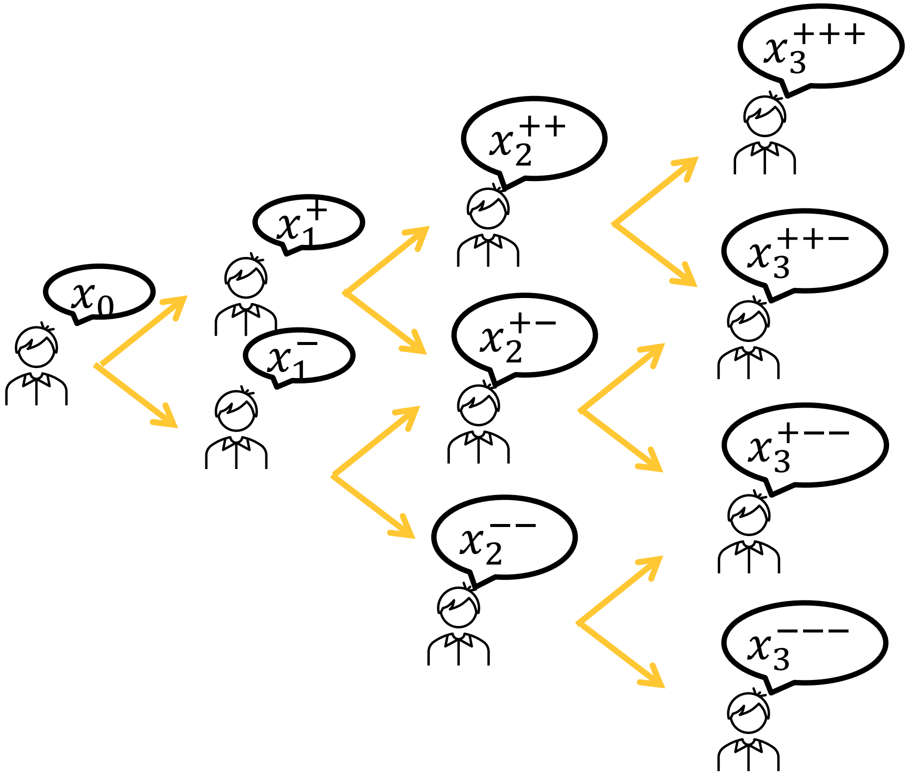

Optimal scoring rule with partial knowledge

- Given $\mathcal{P}$ (a collection of P), design bounded S so that maximize the expected gain

$$\max_{S} \min_{P \in \mathcal{P}} \mathbb{E}_P[S(x, w) - S(x_0, w)]$$
$$\text{s.t. } S \text{ is proper and bounded}$$

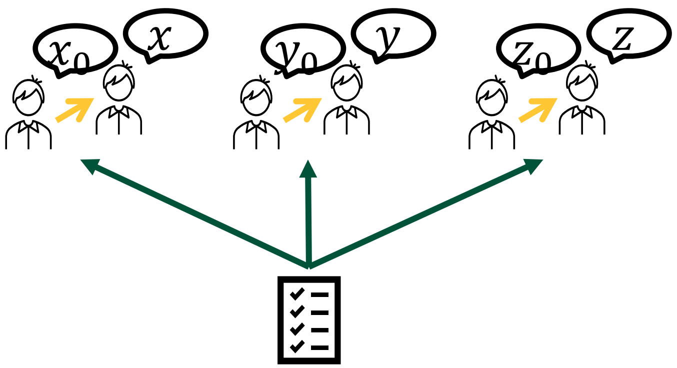

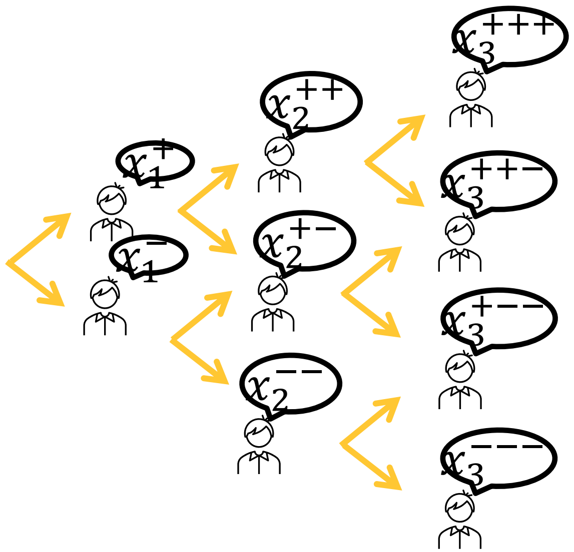

Outline

- Problem set up
  - Proper scoring rule
  - Maximizing information gain
  - Partial knowledge
- Savage characterization
- Main results
  - Core idea: Information gain = convexity
  - Simulation

What are proper scoring rules

- Given a proper scoring rule S
  - Binary outcome: $w \in \{0,1\}$ and prediction: $x \in [0,1]$
  - $S(x, x) \geq S(x', x)$ for all x'.

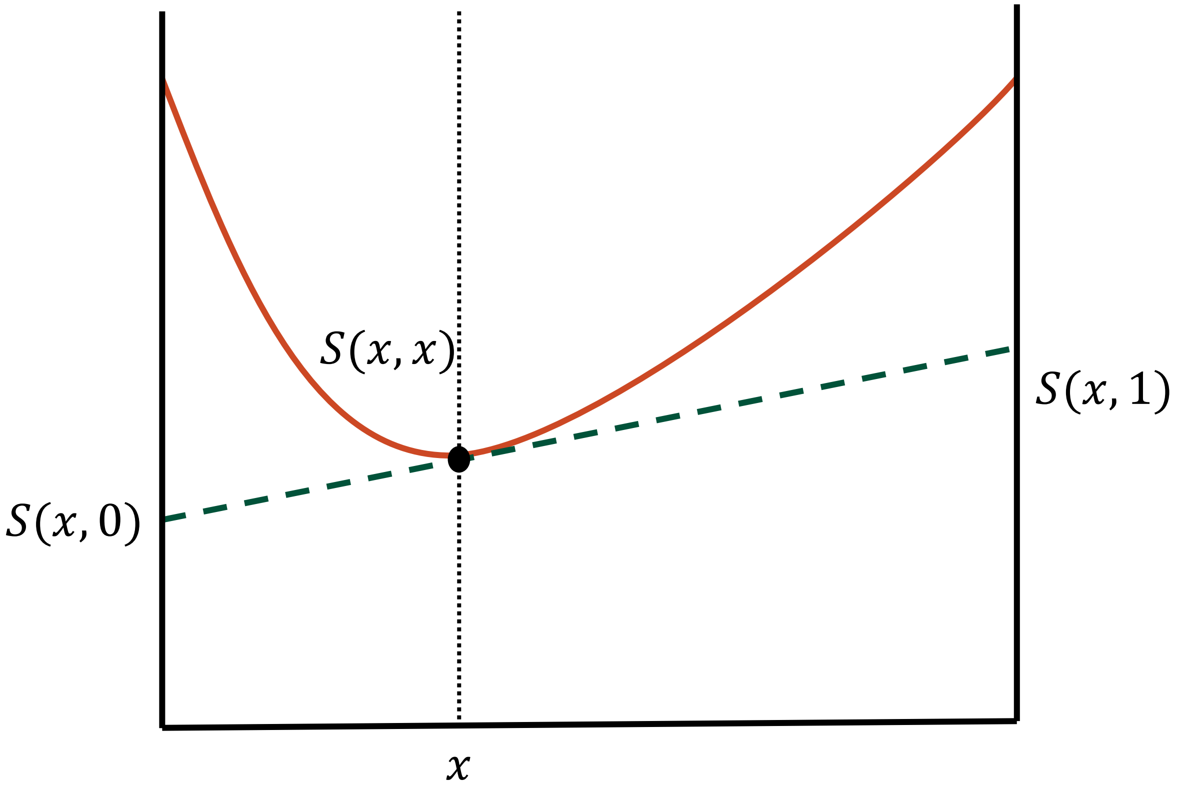

What are proper scoring rules

- Given a proper scoring rule S
  - Binary outcome: $w \in \{0,1\}$ and prediction: $x \in [0,1]$
  - $S(x,x) \geq S(x',x)$ for all x'.

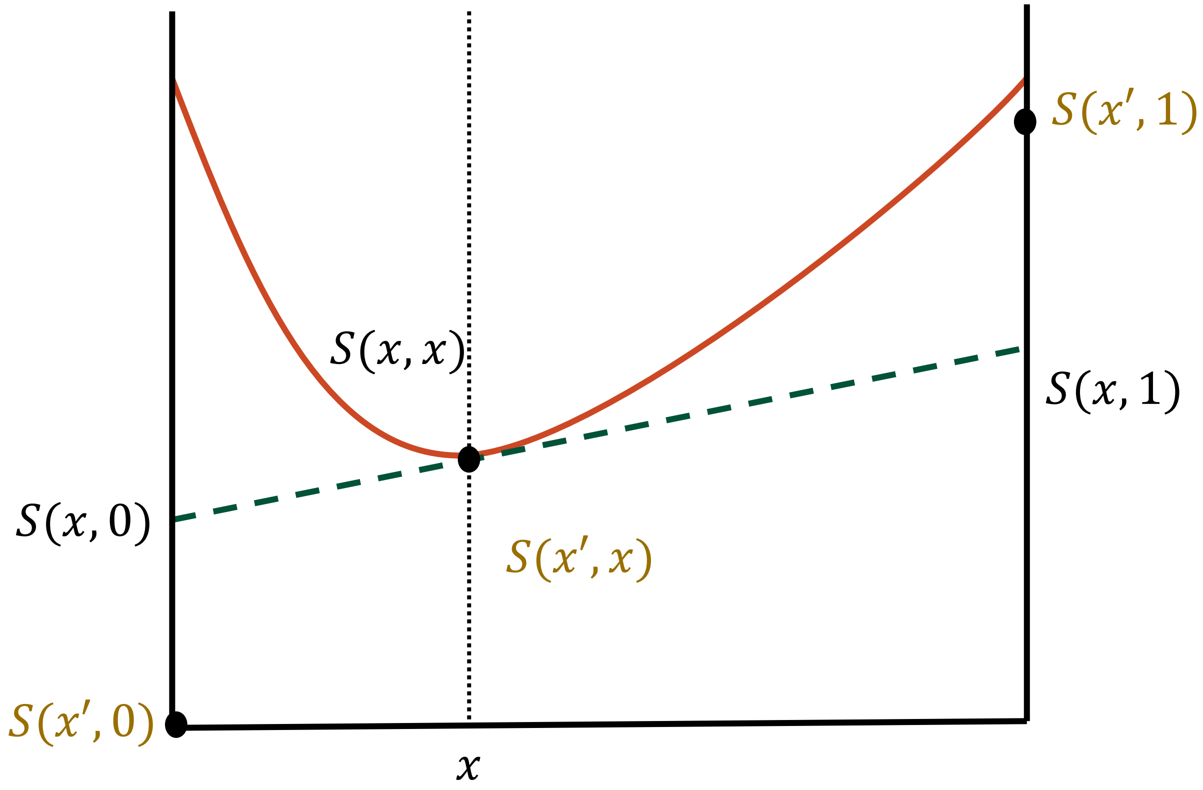

What are proper scoring rules

- Given a proper scoring rule S
  - Binary outcome: $w \in \{0,1\}$ and prediction: $x \in [0,1]$
  - $S(x,x) \geq S(x',x)$ for all x'.

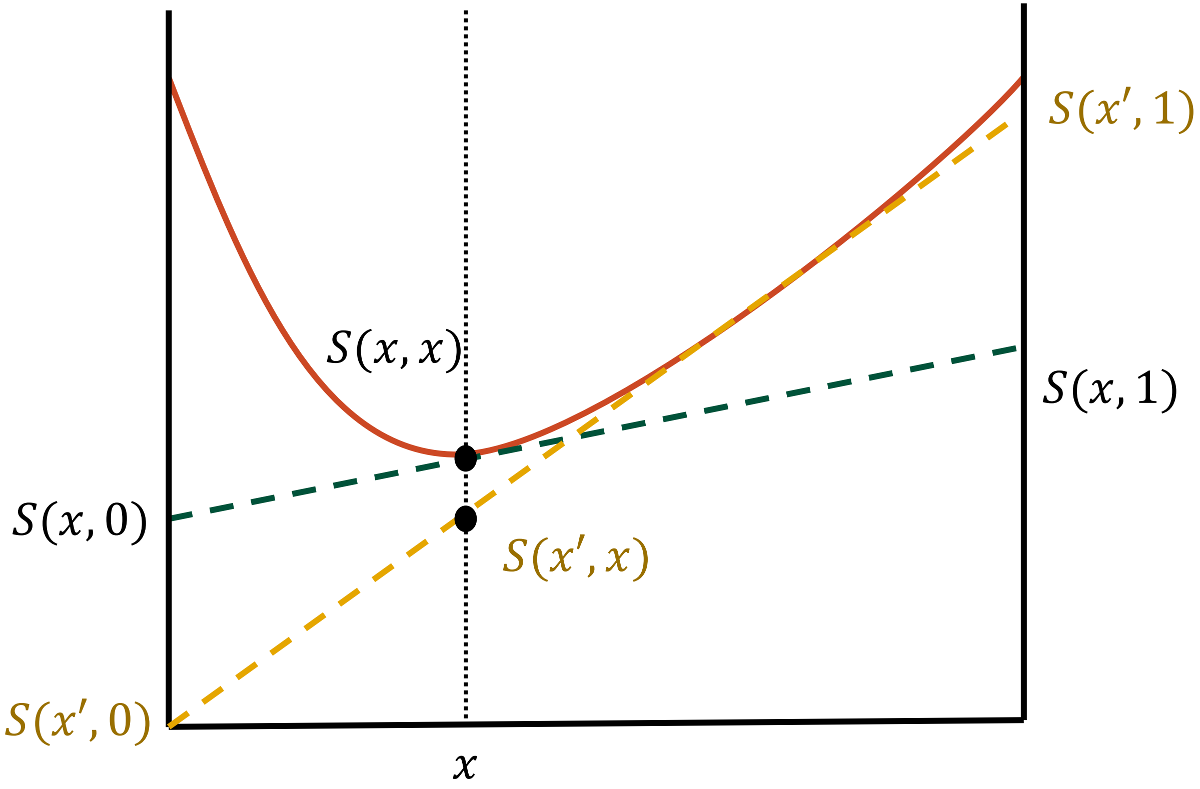

What are proper scoring rules

- Given a proper scoring rule S
  - Binary outcome: $w \in \{0,1\}$ and prediction: $x \in [0,1]$
  - $S(x,x) \geq S(x',x)$ for all x'.

Proper scoring rule S

$\updownarrow$

convex function $H(x) = S(x,x)$

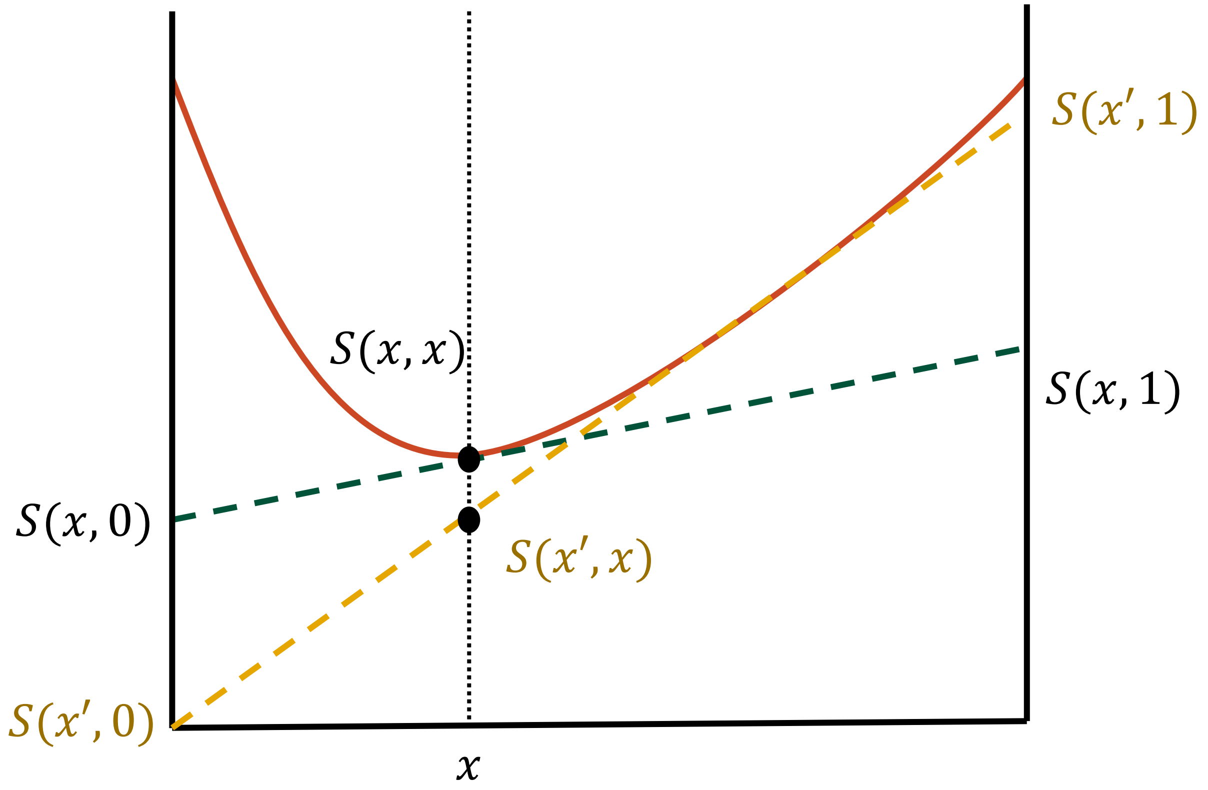

Outline

- Problem set up
  - Proper scoring rule
  - Maximizing information gain
  - Partial knowledge
- Savage characterization
- Main results
  - Core idea: Information gain = convexity
  - Simulation

Main results

- Different $\mathcal{P}$ leads different optimal scoring rules
  - Singleton $\mathcal{P}$: A V-shaped H is optimal.
  - Finite $\mathcal{P}$: An efficient algorithm yields an optimal piecewise linear H
  - "Discretizable" $\mathcal{P}$: An FPTAS for optimal H
    - Homogeneous experiment
    - Beta-Bernoulli

Singleton $\mathcal{P}$

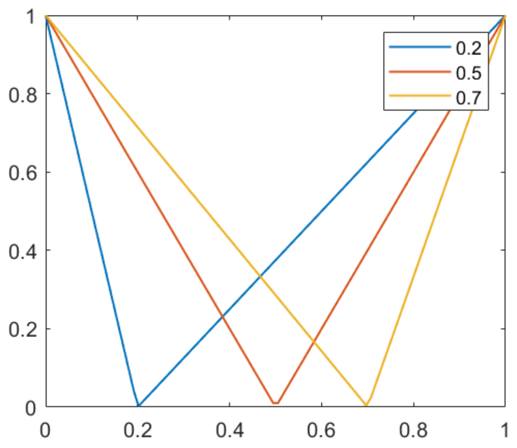

Finite $\mathcal{P}$

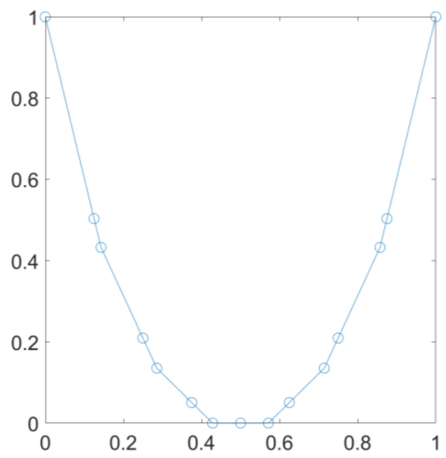

Information gain ↔ convexity

- Example
  - Uniformly distributed $w \sim \{0,1\}$
  - Costly binary signal equals w w.p. $4/5$
  - Prior: $x_0 = 0.5$
  - Posterior:
    - $\Pr[x = 0.8] = \Pr[x = 0.2] = 1/2$

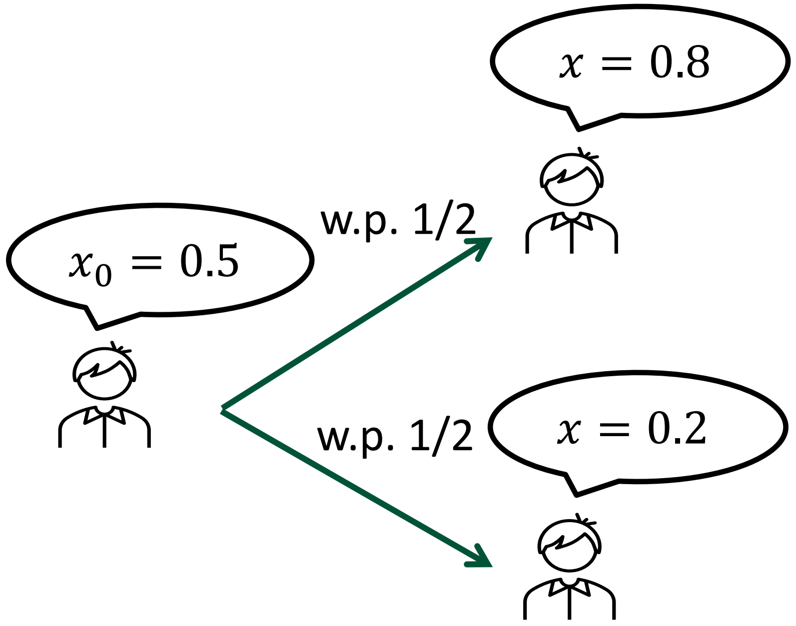

Information gain ↔ convexity

- Example
  - Uniformly distributed $w \sim \{0,1\}$
  - Costly binary signal equals w w.p. $4/5$
  - Prior: $x_0 = 0.5$
  - Posterior:
    - $\Pr[x = 0.8] = \Pr[x = 0.2] = 1/2$

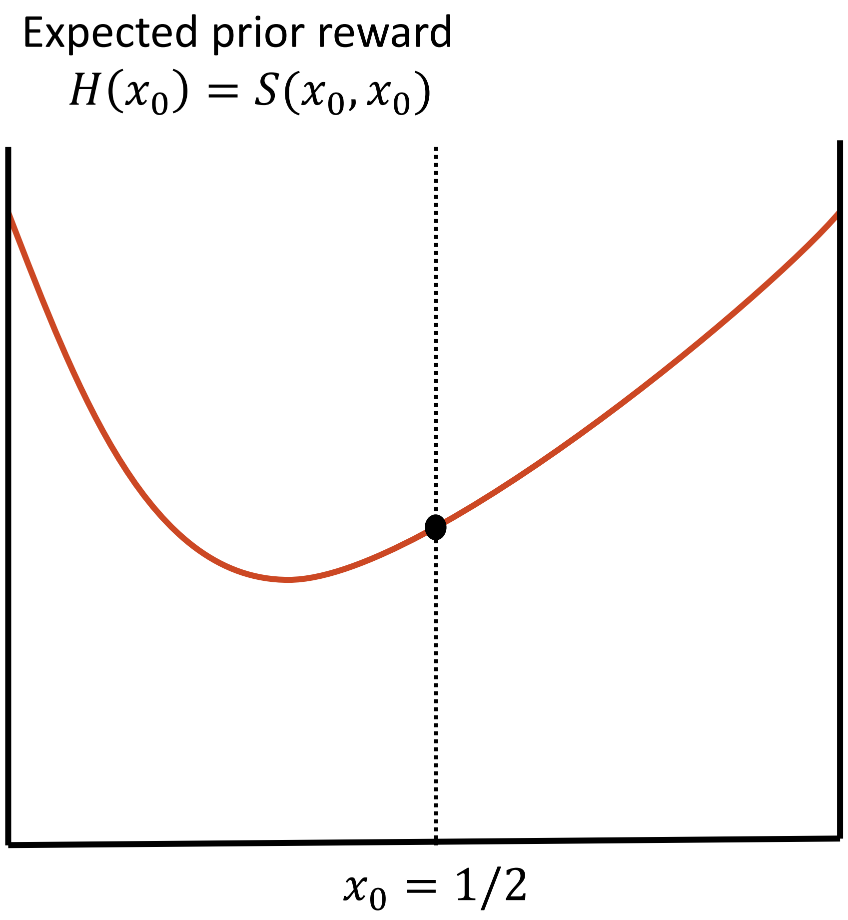

Information gain ↔ convexity

- Example
  - Uniformly distributed $w \sim \{0,1\}$
  - Costly binary signal equals w w.p. $4/5$
  - Prior: $x_0 = 0.5$
  - Posterior:
    - $\Pr[x = 0.8] = \Pr[x = 0.2] = 1/2$
  - Information gain under H
    $\mathbb{E}_P[S(x,w) - S(x_0,w)]$
    $= \mathbb{E}_P[H(x) - H(x_0)]$

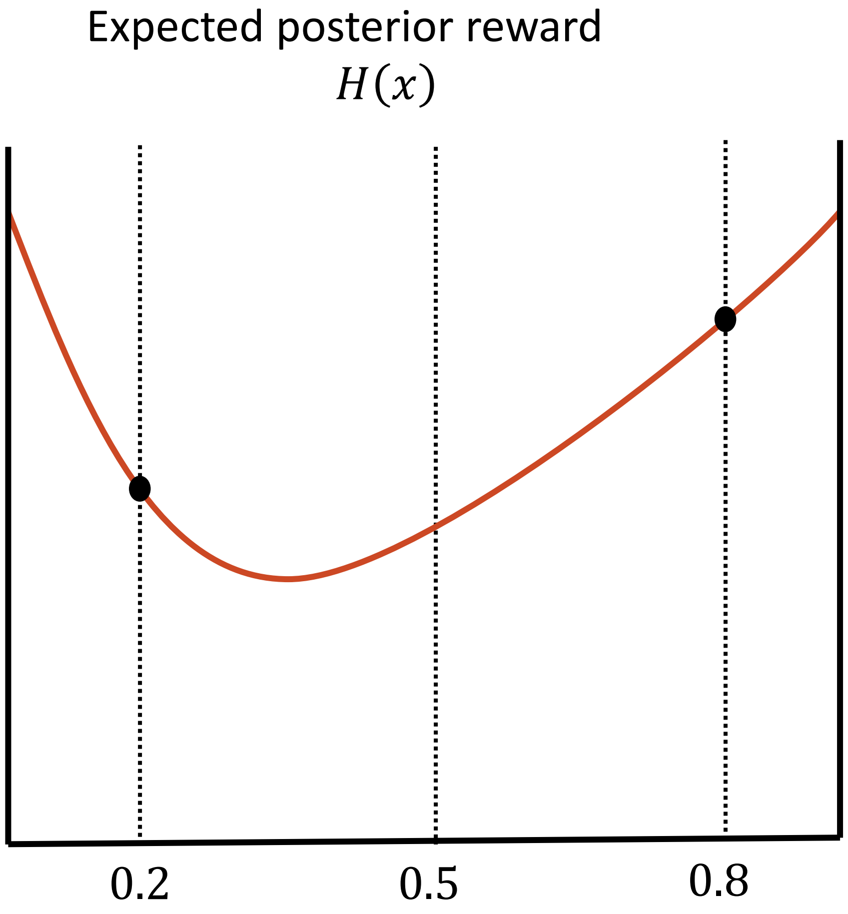

Information gain ↔ convexity

- Example
  - Uniformly distributed $w \sim \{0,1\}$
  - Costly binary signal equals w w.p. $4/5$
  - Prior: $x_0 = 0.5$
  - Posterior:
    - $\Pr[x = 0.8] = \Pr[x = 0.2] = 1/2$
  - Information gain under H
    $\mathbb{E}_P[S(x,w) - S(x_0,w)]$
    $= \mathbb{E}_P[H(x) - H(x_0)]$

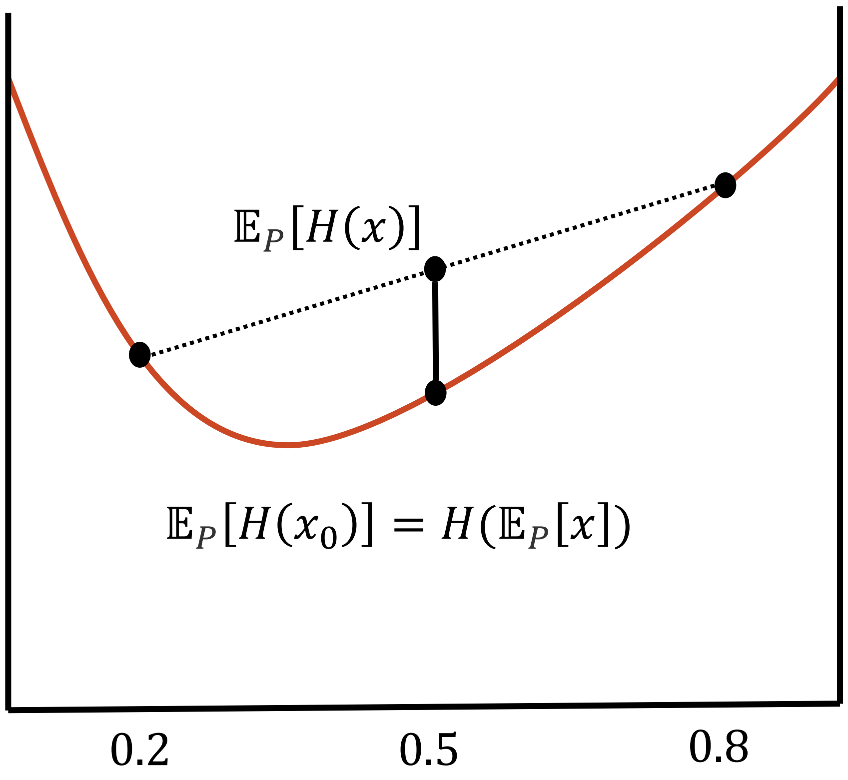

Information gain ↔ convexity

- Example
  - Uniformly distributed $w \sim \{0,1\}$
  - Costly binary signal equals w w.p. $4/5$
  - Prior: $x_0 = 0.5$
  - Posterior:
    - $\Pr[x = 0.8] = \Pr[x = 0.2] = 1/2$
  - Information gain under H
    $\mathbb{E}_P[S(x, w) - S(x_0, w)]$
    $= \mathbb{E}_P[H(x) - H(x_0)]$

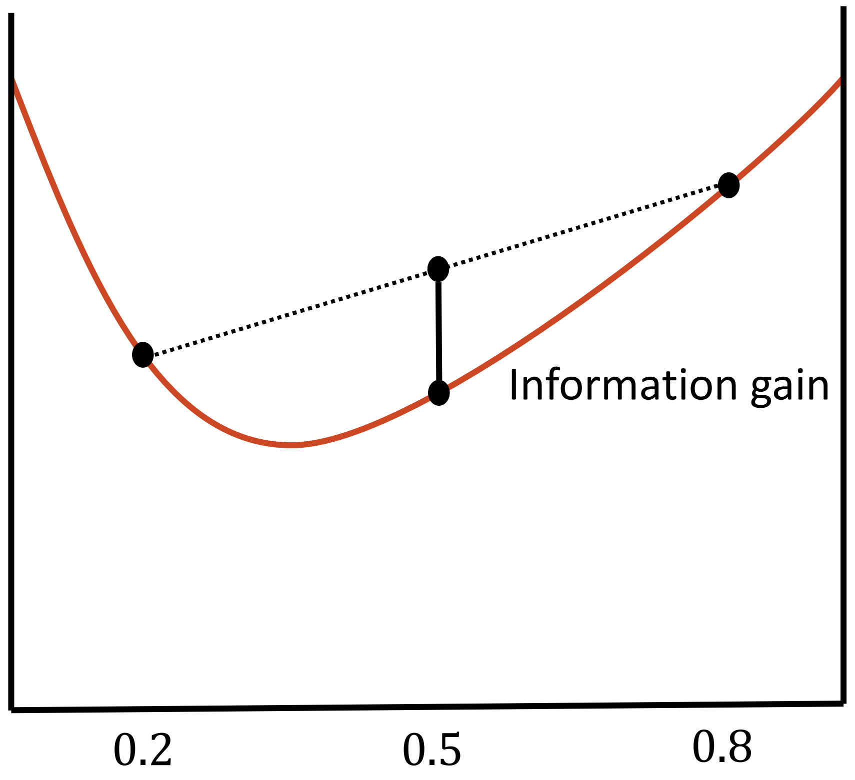

Main results

- Different $\mathcal{P}$ leads different optimal scoring rules
  - Singleton $\mathcal{P}$: a v-shaped H is optimal $\longleftarrow$turning point at prior
  - Finite $\mathcal{P}$: an efficient algorithm and is piecewise linear is optimal $\longleftarrow$turning points at support of $\mathcal{P}$

Singleton $\mathcal{P}$

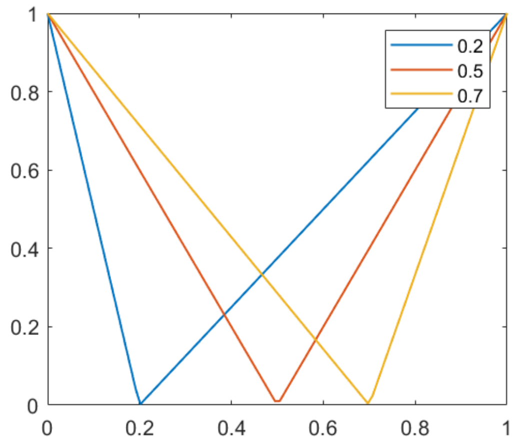

V-shape

Finite $\mathcal{P}$

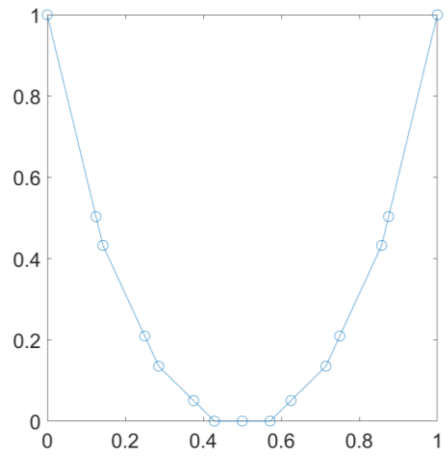

Piecewise linear

Simulations

- Log scoring rule perform well under Beta-Bernoulli setting

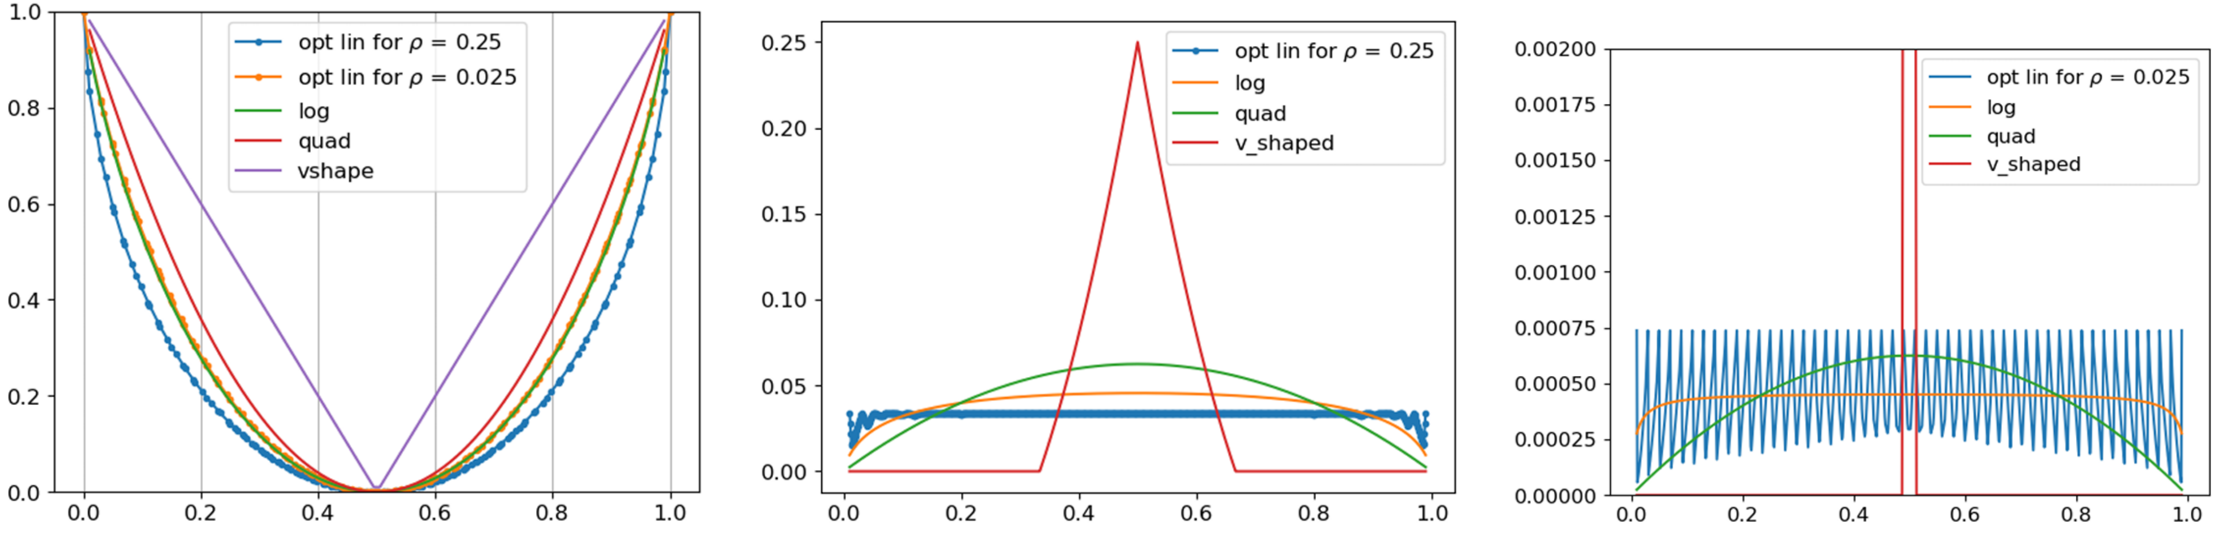

(a) Associated convex functions (b) Information gain with $\rho = 0.25$. (c) Information gain with $\rho = 0.025$.

Simulations

- Log scoring rule perform well under Beta-Bernoulli setting