Scoring rules

- A different kind of mechanism design problem: how to elicit a good prediction of an uncertain event?
  - Weather forecaster: will it rain tomorrow?
  - Political pundit: will a Democrat or Republican win next election?
  - Microsoft employee: will the next version of MS Office ship on time?

Scoring rules

- A different kind of mechanism design problem: how to elicit a good prediction of an uncertain event?
  - Weather forecaster: will it rain tomorrow?
  - Political pundit: will a Democrat or Republican win next election
  - Microsoft employee: will the next version of MS Office ship on time?
- How should we evaluate the quality of a prediction/pay based on the quality of predictions/ incentivize the work needed to output the best possible prediction?

Scoring rules

- X finite set of possible outcomes of uncertain event.
- A scoring rule is a real-valued function $\boxed{S(\vec{q},i)}$
  - $\vec{q}$ is a probability distribution over X (a prediction)
  - i is some outcome in X (the realized outcome)

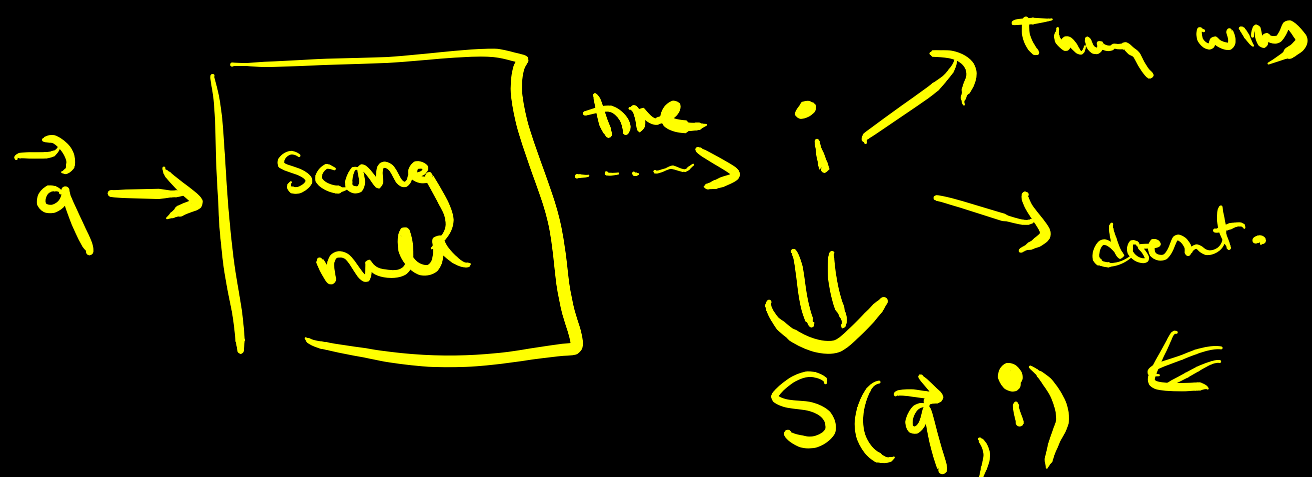

Model for incentives

$X = \{\text{sun}, \text{rain}, \text{snow}\}$

$\quad\;\; p_1, \;\; p_2, \;\; p_3$

$i \in X \qquad p_1 + p_2 + p_3 = 1$

- Forecaster has a belief $\vec{p}$, prob distribution over X.
- Forecaster will choose prediction $\vec{q}$ to maximize expected score

forecaster's goal:

choose $\vec{q}$

to maximize $E_{i \sim p}[S(\vec{q}, i)]$

report $\vec{q}$ $\Rightarrow$ $\left(\sum_{i=1}^{n} p_i \, S(\vec{q}, i)\right)$

Strictly proper scoring rules

- X finite set of possible outcomes of uncertain event.
- A scoring rule is a real-valued function S(q,i)
  - q is a probability distribution over X (a prediction)
  - i is some outcome in X (the realized outcome)

- A scoring rule is strictly proper if, no matter what the true belief $\vec{p}$ of the forecaster is, her unique best response is to report truthfully, i.e. to set $\vec{q} = \vec{p}$.

Strictly proper scoring rules

- X finite set of possible outcomes of uncertain event.
- A scoring rule is a real-valued function S(q,i)
  - q is a probability distribution over X (a prediction)
  - i is some outcome in X (the realized outcome)

$\sum_{i=1}^{n} p_i \, S(\vec{q}, i)$

- A scoring rule is strictly proper if, no matter what the true belief p of the forecaster is, her unique best response is to report truthfully, i.e. to set q = p.

Example: $S(\vec{q}, i) = q_i$

belief $\boxed{(p, 1-p)}$

report $(q, 1-q)$

Exp payoff $= p\,q + (1-p)(1-q)$

given p what q maximizes this?

sd $p = 0.7$

$\boxed{.7q + .3(1-q)}$

Quadratic scoring rule

$$S(\vec{q}, i) = q_i - \frac{1}{2} \sum_{j \in X} q_j^2$$

$q_i = 1$ for some i — if i happens, $1 - \frac{1}{2} = \frac{1}{2}$

$q_j = 0$ $\forall j \neq i$ — if i doesn't happen $-\frac{1}{2}$

$q_i = \frac{1}{n}$ no matter what payoff $\geq \frac{1}{2n}$

$$S(\vec{q}, i) = q_i - \frac{1}{2}\sum_{j \in X} q_j^2$$

GSR is strictly proper

$$E(S_{\text{con}}) = \sum_i p_i q_i - \frac{1}{2}\sum_{i \in X} p_i \sum_{j \in X} q_j^2$$

is max at $\vec{p} = \vec{q}$

$$\frac{d}{dq_k} E(S_{\text{con}}) = p_k - \sum_{i \in X} p_i q_k$$

$$q_k = \frac{p_k}{\underbrace{\sum_{i \in X} p_i}_{=1}} = p_k$$

$aS + b$

$+\frac{1}{2}$

Logarithmic scoring rule

$S(q, i) = \ln q_i$

$|X| = n$

add $\ln |X|$

forecaster can guarantee nonneg exp utility.

$\vec{q} = (\frac{1}{n}, \ldots \frac{1}{n})$

$$E(\text{score}) = \underbrace{\sum_{i=1}^{n} p_i \ln(\tfrac{1}{n})}_{-\ln(n)} + \ln(n) = 0$$

$q_i = 0$ | $q_i \geq \varepsilon$

Logarithmic scoring rule is strictly proper.

- incentivizing honest feedback
- prediction markets

Incentivizing honest feedback

- Example: peer grading, where students grade the assignments of other students.

- How to incentivize accurate grading, without direct verification?

Model

- n players (graders of an assignment, say in MOOC)
- Player i has a "signal" s_i → grader's true opinion of quality of homework.
- Each player submits a report r_i to a mechanism.
- Mechanism pays player $\boxed{\pi_i(r_1, \ldots, r_n)}$

Assume signals $(s_1, \ldots, s_n)$ drawn from distn. $\mathcal{D}$

$n=2$

		bad $s_2=0$	good $s_2=1$
bad	$s_1=0$	0.3	0.1
good	$s_1=1$	0.1	0.5
		0.4	0.6

$Pr(s_2=0 \mid s_1=0) = \frac{3}{4}$

$Pr(s_2=1 \mid s_1=0) = \frac{1}{4}$

How to choose payment fns $\Pi_1(\vec{r}), \ldots \Pi_n(\vec{r})$ to incentivize truthful reporting?

Output Agreement

reward agreement.

- For each player i
  - Pick a random player $j \neq i$
  - Set payoff π_i equal to 1 if they agree, 0 otherwise.

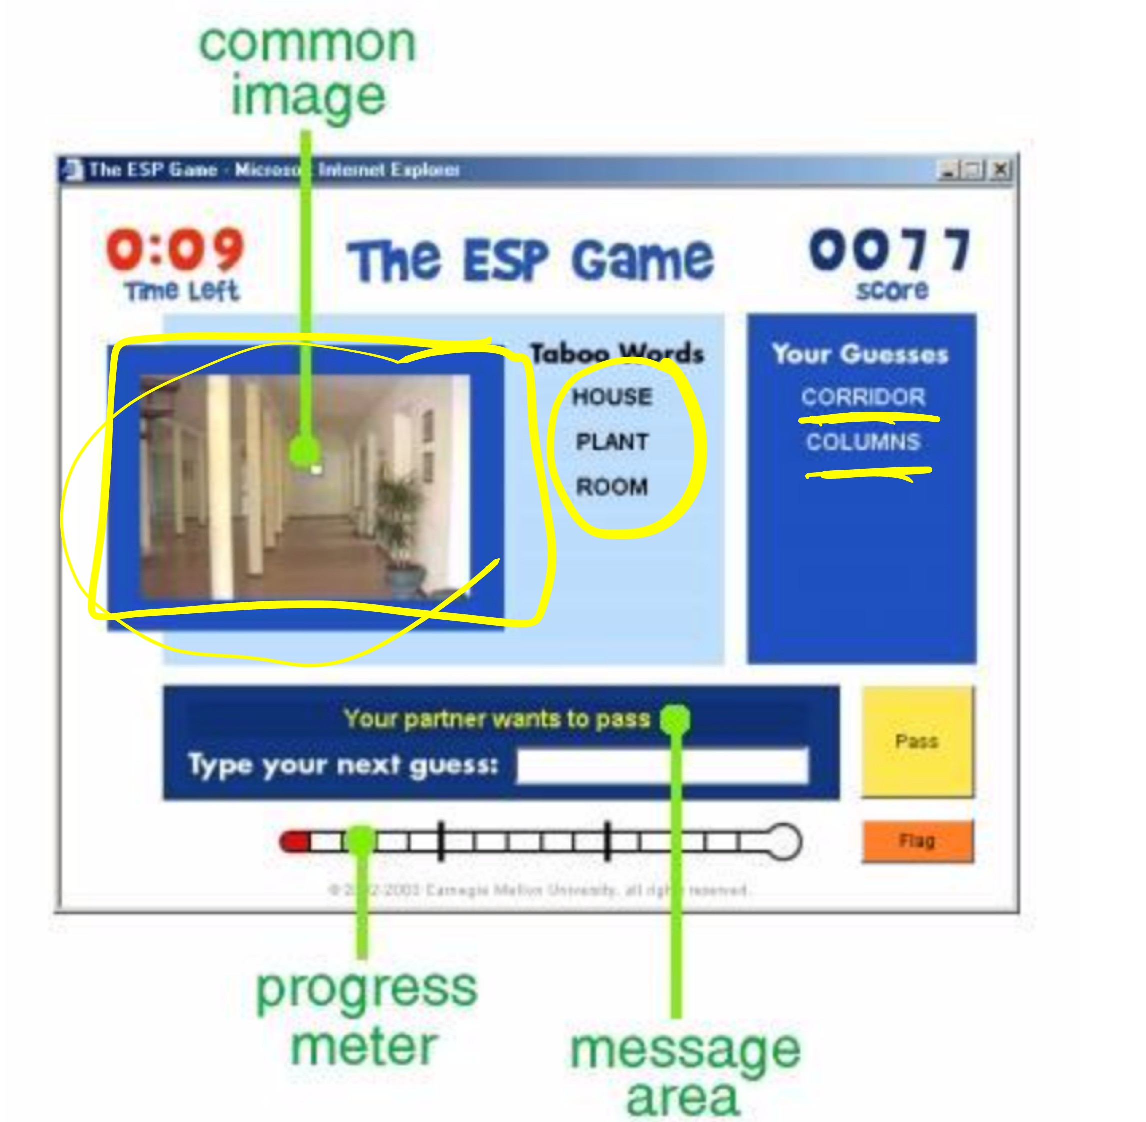
common image
The ESP Game - Microsoft Internet Explorer
0:09
Time Left
The ESP Game
0077
score
Taboo Words
HOUSE
PLANT
ROOM
Your Guesses
CORRIDOR
COLUMNS
Your partner wants to pass
Type your next guess:
Pass
Flag
progress meter
message area

Output Agreement

- For each player i
  - Pick a random player $j \neq i$
  - Set payoff π_i equal to 1 if they agree, 0 otherwise.

Is it a Nash eq to report truthfully

		bad $s_2=0$	good $s_2=1$
bad	$s_1=0$	0.3	0.1
good	$s_1=1$	0.1	0.5
		0.4	0.6

$s_2=0$

$Pr(s_1=0 \mid s_2=0) = \frac{3}{4}$

	0	
	0.1	0.2
1	0.2	0.5

Mechanism has bad NE:

Peer prediction mechanism

- Suppose the distribution D over signals is known to mechanism.

  ? assumptions

- For each player i
  - Pick a random player $j \neq i$
  - Let $D_j(r_i)$ be the distribution of s_j conditioned on $s_i = r_i$
  - Set i's payoff $\pi_i := S(D_j(r_i), r_j)$

$r_1 = 0$

		bad $s_2 = 0$	good $s_2 = 1$
bad	$s_1 = 0$	0.3	0.1
good	$s_1 = 1$	0.1	0.5

$D_2(s_1 = 0) = \left(\frac{3}{4}, \frac{1}{4}\right)$

0 1

$D_2(s_1 = 1) = \left(\frac{1}{6}, \frac{5}{6}\right)$

$S(D_2, 0)$

$S(D_2, 1)$