Computation and Economics - Spring 2012
Section: Social Choice, Reputation Systems and Truthful Elicitation

Professor Sven Seuken
Department of Informatics, University of Zurich

1 Social Choice

Consider the ordinal preferences

$$P = \begin{cases} 4: & a \succ b \succ c, \\ 4: & b \succ c \succ a, \\ 2: & c \succ a \succ b. \end{cases}$$

1.1 Dodgson Rule

The *Dodgson Rule* selects the alternative that can be made a Condorcete winner with the least amount of swaps of adjacent alternatives in the preference orderings of agents. Changing the ranking

$$a \succ b \succ c$$

to

$$a \succ' c \succ' b$$

would be one swap of the adjacent alternatives b and c. However, changing it to

$$c \succ'' b \succ'' a$$

would require 3 swaps of adjacent alternatives, e.g. $a \leftrightarrow b, a \leftrightarrow c, c \leftrightarrow b$.

1.2 Exercise

Apply the Dodgson Rule to P.

- a already wins against b, but looses to c. We can swap c and a once in the last profile, then a and c are tied, which is by definition enough. Thus one swap will make it a Condorcete winner.
- b already wins against c and looses to a. To reach a tie, we could swap a and b in profile 1.
- c already wins against a, but looses to b. To make it a Condocrete winner, we could swap b and c in profiles 1,2,3, which would require 3 swaps.

Hence either a or b are selected.

1.3 Kemeny Rule

The *Kemeny Rule* looks at all possible rank orderings and computes the sum of pairwise supports for each. If a ranking implies $a \succ b$, than the number of rankings implying the same ordering for a and b (and not the opposite) is added to the sum. This is done for all combinations of alternatives. The ranking with the largest support is chosen.

1.4 Exercise

Apply the Kemeny Rule to P.

The values in the pairwise comparison graph are as follows (*from* row *to* column):

	a	b	c
a	-	6	4
b	4	-	8
c	6	2	-

The table below lists all possible orderings and their pairwise support:

Ordering	[pos. 1≻ pos. 2]	pos. [2≻ pos. 3]	[pos. 1≻ pos. 3]	$\sum$
$a \succ b \succ c$	6	8	4	18
$a \succ c \succ b$	4	2	6	12
$b \succ a \succ c$	4	4	8	16
$b \succ c \succ a$	8	6	4	18
$a \succ a \succ b$	6	6	2	14
$c \succ b \succ a$	2	4	6	12

Hence a or b are selected.

1.5 Schulze Rule

The *Schulze Rule* considers paths from each alternative to all others. The path width from a to z is determined as follows:

- Construct the weighted majority graph (WMG) from the pairwise comparison graph by taking differences
- The width of an edge is this number it carries in the WMG, crossing the edge in the opposite direction has the opposite (=negative) weight
- Consider all paths starting in a and ending in z
- Take as the width of each path the minimal width of an edge on that path, which may be negative
- Of all paths from a to z select the one with the greatest width and take this width to be $S(a, z)$

An alternative a is a *Schulze winner* if for all other alternatives z the value $S(a, z) \geq S(z, a)$.

1.6 Exercise

Apply the Schulze Rule to P.

The values in the weighted majority graph are as follows (*from* row *to* column):

	a	b	c
a	-	2	-2
b	-2	-	3
c	2	-3	-

Then the widest paths from each alternative to the others have the following values:

S	a	b	c
a	-	2	2
b	2	-	6
c	2	2	-

We now look for alternatives a with the property: Any other alternative z can be reached via a path that is at least as wide as any path leading from z back to a. This is case for a and b, but not for c.

2 Peer Prediction

Consider the following peer prediction setting, where the quality of a TV set shall be assessed by two peers:

- The TV set has either a high quality (H) or a low quality (L)
- By examining the TV set, each agent reveives a private signal $S_i, i = 1, 2$ in $\{l, h\}$
- The prior is

  $$P(L) = 0.2, P(H) = 0.8,$$

  and the signal distribution is

  $$P(h|H) = 0.8, P(l|L) = 0.9,$$

  and it is common knowledge to the agents and the mechanism
- The mechanism collects a report $\hat{s}_i, i = 1, 2$ from the agents
- The mechanism uses the reports to compute $g(\hat{s}_j|\hat{s}_i) = P(S_j = \hat{s}_j|S_i = \hat{s}_i), i \neq j$, the conditional probability of j observing signal $\hat{s}_j$, given i observes signal $\hat{s}_i$
- Agent i is scored using the logarithmic scoring rule, i.e. $R(\hat{s}_i, \hat{s}_j) := \ln(g(\hat{s}_j|\hat{s}_i)), i \neq j$.

What are the conditional probabilities that agent 2 receives signal h, i.e. $g(h|h) = P(S_2 = h|S_1 = h)$ and $g(h|l) = P(S_2 = h|S_1 = l)$?

First note that[1]

$$g(h|h) = P(S_2 = h|H)P(H|S_1 = h) + P(S_2 = h|L)P(L|S_1 = h).$$

[1] We use conditional independence of the signals S_1 and S_2 when the quality is fixed (H or L).

We know $P(S_2 = h|H)$ and $P(S_2 = h|L)$, but not the other terms. So we write

$$\begin{aligned} P(H|S_1 = h) &= \frac{P(S_1 = h|H)P(H)}{P(S_1 = h)} \\ &= \frac{P(S_1 = h|H)P(H)}{P(S_1 = h|H)P(H) + P(S_1 = h|L)P(L)} \\ &= \frac{0.8 \cdot 0.8}{0.8 \cdot 0.8 + 0.2 \cdot 0.2} = 0.94. \end{aligned}$$

Consequently, $P(L|S_1 = h) = 1 - P(H|S_1 = h) = 0,06$. Thus, $g(h|h) = 0.8 \cdot 0.94 + 0.2 \cdot 0.06 = 0.76$.

Similarly, we get

$$g(h|l) = P(S_2 = h|H)P(H|S_1 = l) + P(S_2 = h|L)P(L|S_1 = l),$$

$$\begin{aligned} P(H|S_1 = l) &= \frac{P(S_1 = l|H)P(H)}{P(S_1 = l)} \\ &= \frac{P(S_1 = l|H)P(H)}{P(S_1 = l|H)P(H) + P(S_1 = l|L)P(L)} \\ &= \frac{0.2 \cdot 0.8}{0.2 \cdot 0.8 + 0.9 \cdot 0.2} = 0.47, \end{aligned}$$

and $g(h|l) = 0.53$.

Assume that agent 1 gets signal $S_1 = h$ and that agent 2 will report $\hat{s}_2 = S_2$ truthfully. Calculate the expected scores for agent 1 from reporting truthfully and from lying.

The score from reporting truthfully depends on the signal reported by agent 2, i.e.

$$\begin{aligned} E^T(R) &= g(h|h)R(h,h) + g(l|h)R(h,l) \\ &= g(h|h)\ln(g(h|h)) + (1 - g(h|h)\ln(g(l|h)) \\ &= 0.76\ln(0.76) + (1 - 0.76)\ln(1 - 0.76) = -0.55. \end{aligned}$$

The score from reporting l is

$$\begin{aligned} E^L(R) &= g(h|h)R(l,h) + g(l|h)R(l,l) \\ &= g(h|h)\ln(g(h|l)) + (1 - g(l|h)\ln(g(l|l)) \\ &= 0.76\ln(0.53) + (1 - 0.76)\ln(1 - 0.53) = -0.66. \end{aligned}$$

3 Proper Scoring Rules

A scoring rule is a function of the form

$$R : \Omega \times \mathcal{P} \to \mathbb{R} \cup \{+\infty\}.$$

It takes as inputs:

- A probability distribution $\hat{P} \in \mathcal{P}$ on Ω, which the agent reports to the mechanism
- An observed outcome $\omega \in \Omega$

The it assigns a score (reward/payoff/utility), based on these inputs. The scoring rule is *proper* if the agent maximizes his expected payoff by reporting his true beliefs about the distribution.

3.1 Exercise

Suppose a setting with binary outcome, e.g. H or L. The agent believes that $p = 0.1$ is the true probability of H. He reports a (potentially different) probability $\hat{p}$ to the mechanism. The mechanism scores the agent by paying him the probability he predicted for the event that actually occured, i.e. $R(\omega, (\hat{p}, 1 - \hat{p})) = \hat{p}\mathbf{1}_{\{\omega=H\}} + (1 - \hat{p})\mathbf{1}_{\{\omega=L\}}$.

What is the expected payoff to the agent as a function of the report $\hat{p}$? What should the agent report to maximize the expected payoff?

His expected payoff (before the actual outcome is known) is

$$E^{(p,1-p)}[R(\cdot, (\hat{p}, 1 - \hat{p})] = R(H, (\hat{p}, 1 - \hat{p}))p + R(L, (\hat{p}, 1 - \hat{p}))(1 - p).$$

Since $\hat{p}$ is the only thing the agent can manipulate, we can interprete this as a function where the report is mapped to an expected payoff:

$$\hat{p} \mapsto E^{(p,1-p)}[R(\cdot, (\hat{p}, 1 - \hat{p})].$$

Then the expected payoff depends on the report $\hat{p}$ and is

$$E^{(p,1-p)}[R(\cdot, (\hat{p}, 1 - \hat{p}))] = p\hat{p} + (1 - p)(1 - \hat{p}) = -0.8\hat{p} + 0.9.$$

Of course this scoring rule is not proper, since the agent maximizes this value by reporting $\hat{p} = 0 \neq 0.1$.