CSCI699: Topics in Learning & Game Theory

Lecture 10

Lecturer: Shaddin Dughmi *Scribes: Omkar Thakoor*

Crowdsourcing Information

In this setting, we have a *principal*, and n *agents*. We have a set Ω of possible *states of nature*, where a particular $\omega \in \Omega$ is drawn from a common prior $P(\omega)$. Each agent i receives a signal (a.k.a. type) t_i from a finite set T of possible types. For any agent i, and any $t \in T$, $P(t|\omega)$ denotes $Pr[t_i = t|\omega]$, thereby, $t_1, \ldots, t_n$ are conditionally i.i.d. given any $\omega \in \Omega$. The goal of the principal is to incentivize agents to report their types truthfully, in order to update his belief about ω. That is, he wants to learn the posterior $P(\omega|t_1, \ldots, t_n)$.

Running Example

Let the state of nature denote whether a new iPhone edition is good, or bad. Accordingly, let $\Omega = \{G, B\}$ with the associated prior $P(G) = 0.6$, $P(B) = 0.4$. Let the types of people denote whether they like or dislike the iPhone, thus, let $T = \{L, D\}$. When the iPhone is good, let the type of an agent be distributed as per $P(L|G) = 0.75$ ($\Leftrightarrow P(D|G) = 0.25$) and, similarly, let the distribution when the iPhone is bad, be $P(D|B) = 0.75$ ($\Leftrightarrow P(L|B) = 0.25$).

In this case, if $n = 4$ agents report their types to be L, L, L, D respectively, the principal's posterior estimate is $P(G|L, L, L, D) = \frac{0.6 \cdot 0.75^3 \cdot 0.25}{0.6 \cdot 0.75^3 \cdot 0.25 + 0.4 \cdot 0.25^3 \cdot 0.75} \approx 0.93$ ($\Leftrightarrow P(B|L, L, L, D) \approx 0.07$)

This problem can be modeled in several ways.

Model 1: Observable Model

In this model, we make the following assumptions.

We assume that the principal solicits the types from agents today, ω is directly revealed to the principal tomorrow, and the principal may pay to the agents day after tomorrow.

We also assume that the principal knows $P(\omega)$ and $P(t|\omega)$. (Equivalently, he knows the joint distribution $P(\omega, t)$.)

Finally, we make a technical assumption that types $t \neq t'$ induce distinct posterior distributions w.l.o.g., i.e.$\exists\ \omega \in \Omega : P(\omega|t) = P(\omega|t')$.

With the last assumption, reporting a type t can be considered equivalent to reporting the distribution $P(\omega|t)$. Hence, the principal can simply use a strictly proper scoring rule to incentivize the agents to report the respective distributions $P(w|t)$ truthfully. With these truthful reports, with the knowledge of $P(\omega)$ and $P(t|\omega)$ as mentioned in the second model assumption, and the inherent property that t_i's are i.i.d. given any ω, the principal can apply the Bayes' theorem to compute his posterior estimate.

Next, we consider another model which may suit for certain problem scenarios.

Model 2: Peer Prediction

This model consists of the following assumptions.

As in the case of the Observable model, we assume that the principal knows the joint distribution $P(\omega, t)$.

A key difference with the previous model is that the principal never directly sees ω. He must decide payments solely based on type reports $r_1, \dots, r_n$.

Finally, we make a technical assumption. Consider $P(t_j|t_i) = \sum_\omega P(\omega|t_i)P(t_j|\omega)$. We assume that $P(t_j|t_i = t) \neq P(t_j|t_i = t')\ \forall t \neq t'$.

To intuitively understand this, consider the iPhone example introduced earlier. One would expect that agent 1 liking iPhone should raise the probability estimate that 2 likes as well. (Note that this holds for the distributions P in general position, even by fixing $P(t|\omega)$.)

Now, as per the last assumption, agent i reporting a type is equivalent to him reporting a posterior distribution on the type of some agent $j \neq i$. Let $q_t \in \Delta(T)$ be the conditional distribution of t_j given $t_i = t$ for any $j \neq i$. (Note that it doesn't depend on i, j.) Then, the *Peer prediction* protocol is as follows:

1. To each agent i, assign agent $\hat{i} \neq i$ as his *peer*.
2. Solicit type reports $r_1, \dots, r_n$ from the agents, with each $r_i \in T$.

3. Let $S : \Delta(T) \times T \to \mathbb{R}$ be a strictly proper scoring rule. Pay agent i a value of $S(q_{r_i}, r_{\hat{i}})$ (computation depends on $P(\omega)$).

For the solution above, we prove the following result.

Theorem 1. *In the Peer prediction protocol, reporting $r_i = t_i \ \forall i$ is a strict Bayes-Nash equilibrium.*

Proof. Assume all agents except i (including $\hat{i}$) report truthfully. Given t_i, i believes $t_{\hat{i}} \sim q_{t_i}$). In step 3 of the protocol, the payment to i is $S(q_{r_i}, t_{\hat{i}})$ (since we assume $r_{\hat{i}} = t_{\hat{i}}$). Since S is a strictly proper scoring rule, i's best strategy is to report posterior q_{t_i} on $t_{\hat{i}}$, hence he must report $r_i = t_i$.

□

Model 3 : Bayesian Truth Serum

The assumptions here are as follows:

The principal does not know $P(\omega)$, nor $P(t|\omega)$ (even though he knows they exist and that they determine the agents' behavior).

We assume $n \to \infty$ (This assumption can be removed by using results from follow-up work).

We propose a protocol, called the *Bayesian truth serum* protocol. The idea is as follows. The principal solicits from agent i, not only the type report $r_i \in T$, but also a prediction $y^i \in \Delta(T)$ of the the empirical type frequency $\bar{x}$ with each $\bar{x}_j$ defined as the fraction of agents reporting j as their type. Having received r_i, and y^i from each agent i, the principal rewards him for a "surprisingly common" type report, i.e., if r_i is more common in the empirical distribution than estimated by the other agents. Secondly, an agent is also rewarded for a truthful prediction y^i using a strictly proper scoring rule. Formally,

1. Solicit type report $r_i \in T$, as well as a prediction$y^i \in \Delta(T)$ of the the empirical type frequency $\bar{x}$.

2. Define $\bar{y}$ so that $\forall k \in T : \log \bar{y}_k = \frac{1}{n} \sum_i \log y^i_k$.

3. Pay agent i a value of $F(r_i) + G(y^i)$, where, $F(j) = \log \frac{\bar{x}_j}{\bar{y}_j}$ is the reward for "surprisingly common" type reports, and, $G(y^i) = E_{k \sim \bar{x}}[S(y^i, k)]$ is the reward for accurately predicting $\bar{x}$ with y^i, (by choosing S to be a strictly proper scoring rule).

$G(y^i)$ can be denoted as $S(y^i, \bar{x})$ as per the notation established in Lecture 9. Clearly, this is maximized at $y^i = \bar{x}$.

In the next lecture, we will prove the following result for this protocol.

Theorem 2. *In the Bayesian truth serum protocol, Each agent i reporting $r_i = t_i$ and y^i as their posterior belief on $\bar{x}$ given t_i, is a Bayes-Nash equilibrium.*