NETS 412: Algorithmic Game Theory	April 20, 2017
Lecture 23	
Lecturer: Bo Waggoner	*Scribe: Bo Waggoner*

Proper Scoring Rules and Prediction Markets

Today we look at the question: How can we incentivize an agent or group of agents to make an accurate prediction? This is actually a simpler question than most of the mechanism-design settings considered in this class so far: our agents won't be getting items, only reporting predictions and receiving payments based on their accuracy.

1 Proper Scoring Rules

There is a future event or random variable Y with a finite set $\mathcal{Y}$ of possible outcomes. For example, $\mathcal{Y} = \{\text{sunny, cloudy}\}$. Let $\Delta_{\mathcal{Y}}$ be the set of probability distributions on $\mathcal{Y}$.

1. A single agent, the "expert", reports a probability distribution $p \in \Delta_{\mathcal{Y}}$. This is interpreted as a prediction of the chance of each outcome.
2. The mechanism then observes the true outcome, say y.
3. The mechanism gives the agent a "score" according to a function $S : \Delta_{\mathcal{Y}} \times \mathcal{Y} \to \mathbf{R}$. Their score is $S(p, y)$.

We can interpret the score as a payment that the mechanism will give to the agent. **We assume that the agent's goal is always to maximize expected score.** Suppose that the agent believes the true distribution is q and she reports p. Then let us use the notation $S(p; q)$ for her expected score:

$$S(p; q) := \mathrm{E}_q S(p, Y) = \sum_y q(y) S(p, y).$$

Definition 1 *A* scoring rule *is a function $S : \Delta_{\mathcal{Y}} \times \mathcal{Y} \to \mathbf{R}$. It is* proper *if truthfulness maximizes expected score: for all beliefs q and all $p \neq q$*

$$S(q; q) \geq S(p; q).$$

It is strictly *proper if truthfulness uniquely maximizes expected score: for all q and $p \neq q$,*

$$S(q; q) > S(p; q).$$

Let δ_y be the distribution putting probability one on outcome y. Notice that if an agent believes δ_y, then she will definitely receive $S(p, y)$ for prediction p. So $S(p; \delta_y) = S(p, y)$.

1.1 Examples

One proper scoring rule is the log scoring rule: $S(p, y) = \log p(y)$.

Let's check that it's strictly proper. The expected score is $S(p; q) = \sum_y q(y) \log p(y)$. The expected score for truthfulness is $S(q; q) = \sum_y q(y) \log q(y)$.[1] So the difference is

$$\begin{aligned} S(q; q) - S(p; q) &= \sum_y q(y) \left(\log q(y) - \log p(y) \right) \\ &= \sum_y q(y) \log \frac{q(y)}{p(y)}. \end{aligned}$$

[1] Where have you seen this expression before? A: It's the negative of the Shannon entropy of q!

This is actually equal to what's called the KL-divergence or relative entropy between q and p, written $KL(q, p)$. It is known to be nonnegative, and strictly positive for $p \neq q$, which proves strict properness. You can also check properness manually with some calculus.[2]

A second example is the quadratic scoring rule $S(p, y) = 2p(y) - \sum_{y' \in \mathcal{Y}} p(y')^2$. Exercise for the reader: Compute $S(p; q)$ and show that it is strictly proper.

1.2 Characterization

A natural next question is: What are *all* proper scoring rules? There turns out to be a nice characterization in terms of *convex functions*.

Recall that a function $G : \mathbf{R}^n \to \mathbf{R}$ is *convex* if the points lying above the function form a convex set, i.e. the set $S = \{(x, y) : y \geq G(x)\}$ is convex. In particular, for any two points on or above G, the line segment between them is also on or above G. The function is *strictly convex* if the line segment always lies strictly above G except for its endpoints.

Recall that the *gradient* of a function G is the multidimensional derivative: $\nabla G(x) = \left(\frac{\partial G}{\partial x_1}(x), \ldots, \frac{\partial G}{\partial x_n}(x)\right)$. This is the slope, along each coordinate, of the plane tangent to G at x. Of course, a gradient may not always exist if the function is not differentiable. However, it turns out that convex functions always have a *subgradient*, which intuitively generalizes the gradient to be the slope of any plane tangent to G at a point. A *subgradient* of a convex function G at x is a vector $dG(x)$ such that, for all y, we have $G(y) \geq G(x) + dG(x) \cdot (y - x)$. Note the geometric interpretation: Starting at location x and height $G(x)$, move from x to y, moving vertically according to the slope $dG(x)$. The height is $G(x) + dG(x) \cdot (y - x)$, which lies below the value of the function $G(y)$. A final note is that G is strictly convex if and only if, for all $x \neq y$ and subgradients $dG(x)$, we have $G(y) > G(x) + dG(x) \cdot (y - x)$.

Theorem 2 *A scoring rule $S : \Delta_{\mathcal{Y}} \times \mathcal{Y} \to \mathbf{R}$ is proper if and only if there exists a convex function $G : \Delta_{\mathcal{Y}} \to \mathbf{R}$ with $G(q) = S(q; q)$ and*

$$S(p, y) = G(p) + dG(p) \cdot (\delta_y - p)$$

where $dG(p)$ is a subgradient of G at p and δ_y is the distribution putting probability 1 on outcome y.

Notice that $G(q)$ is the expected score when having belief q and reporting truthfully.

Proof First, suppose there exists a convex function G with $S(p, y) = G(p) + dG(p) \cdot (\delta_y - p)$. We'll prove that S is proper. The expected score with belief q is

$$\begin{aligned} S(p; q) &= \sum_y q(y) S(p, y) \\ &= G(p) + dG(p) \cdot (q - p) \end{aligned}$$

using linearity of expectation. (If you aren't sure exactly how we got this, then it's worth working it out by using the definition of the dot product.) In particular, $S(q; q) = G(q)$ as desired. Now notice by definition of the subgradient, $S(p; q) \leq G(q) = S(q; q)$, which proves that S is proper.

Second, suppose that S is proper. The key point is that for each p, the function $S(p; \cdot)$ is linear. We can define $G(q) := S(q; q)$. Now notice that for all $p \neq q$, we have $G(q) \geq S(p; q)$. So G lies above each of the linear functions $S(p; \cdot)$ and is equal to $S(p; q)$ at $p = q$. In other words, G is a pointwise maximum of linear functions, so G is a convex function. Furthermore, each $S(p; \cdot)$ is a linear function tangent to G at p, so we can write $S(p; q) = G(p) + dG(p) \cdot (q - p)$ for some subgradient $dG(p)$ of G at p. As discussed above, we have $S(p, y) = S(p; \delta_y) = G(p) + dG(p) \cdot (\delta_y - p)$. ■

[2] Notice an interesting case is when $p(y) = 0$ for some y with $q(y) > 0$. What is the score $S(p, y)$ in this case, and what's its interpretation?

Some more exercises for the reader:

- Check that the above theorem holds if we replace "convex" with "strictly convex" and "proper" with "strictly proper".
- Show that, if $S(p, y)$ is (strictly) proper, then $\alpha S(p, y) + \beta$ is also strictly proper for any positive constant α and constant β.
- Consider the case with two outcomes, and draw (or plot with software) the function $G((r, 1-r)) = r^2 + (1-r)^2$. Compute the associated proper scoring rule S. Plot e.g. $S(0.3; \cdot)$ and $S(0.7; \cdot)$ and check that they act as the proof of the theorem says they should.

2 Aggregating Predictions

What if we wanted to get the collective opinion of a large group of experts? One thing we can do is ask them all to simultaenously report a prediction, and score each one with a proper scoring rule. This does elicit each's belief truthfully, but there are some drawbacks:

- It may be expensive or wasteful: If the experts all agree, then we pay many times for the same prediction/information.
- On the other hand, if the experts give many different predictions, then it isn't clear how to **aggregate** them into a single prediction. Should we just average them? (Can you think of a scenario where this would be a bad idea?)
- In particular, it might require expert knowledge to correctly aggregate the predictions.

Robin Hanson had an excellent idea for addressing both of these in his paper *Combinatorial Information Market Design* (2003).

The idea is to approach the experts *sequentially* and ask them to make a prediction, given the predictions that previous experts have made. However, they are expected to "improve" on the prediction: they will get the score of their prediction *minus* the score of the previous one. If we denote the ith expert's prediction by $p^{(i)}$, then she is scored with $S(p^{(i)}, y) - S(p^{(i-1)}, y)$ when the outcome is y. (The designer can choose some initial prediction $p^{(0)}$ so that the first expert's score is well-defined.)

Hanson called this a *market scoring rule* form of a "prediction market".

Theorem 3 *If S is a proper scoring rule, then each agent maximizes expected score by reporting her true belief given her own knowledge and the trades of the previous experts.*

Proof The ith expert will be scored with $S(p^{(i)}, y) - S(p^{(i-1)}, y)$. The second term is a constant not in her control, so her score is $S(p^{(i)}, y) - \beta$ for some constant β, which is proper. So she maximizes expected score by reporting her true belief at the time she is asked. ■

This assumes crucially that each agent is only approached **once**. If they can participate multiple times, then they may have an incentive to misreport early and only truthfully report later.

2.1 Prediction Markets

Real-world prediction markets work more like stock markets: Agents arrive and can buy or sell "shares" of "securities", where each security is named after a possible outcome in $\mathcal{Y}$ and a share in the security for y pays off 1 if $Y = y$ and 0 otherwise.[3]

[3]There are also more complex securities with more complicated payoff functions, but that's beyond the scope of this lecture.

The idea is that the price of the security should approach a "collective prediction" for the probability of the event, since the probability of the event is exactly the expected value of a share.

It turns out that we can re-interpret the market scoring rule process above to look like a prediction market with securities and shares. For example, suppose we use the quadratic scoring rule with a binary event $\mathcal{Y} = \{\text{rain}, \text{no rain}\}$. So $S(p, y) = 2py - p^2$ where $y = 1$ if it rains and 0 otherwise, with p the probability of rain.

An agent moves the prediction from $p^{(0)} = 0.5$ to $p^{(1)} = 0.6$. If it rains, she'll get $S(0.6, 1) - S(0.5, 1) = 0.09$; if not, she'll get $S(0.6, 0) - S(0.5, 0) = -0.11$. We can interpret these payments as buying 0.2 shares of the "rain" security for a price of 0.11. If it rains, each share pays off one dollar, so she gets 0.2 dollars, minus the 0.11 spent for a profit of 0.09. If it doesn't rain, the fraction of a share pays off nothing and she loses 0.09.

We can broaden this into a general, mathematical connection: a prediction market, given a reasonable rule for how to update the price based on trades, is equivalent in incentives to Hanson's market scoring rule. The formal mathematical connection involves duality of convex functions and is very cool, but beyond the scope of this lecture! The main paper to read is *Efficient Market Making via Convex Optimization, and a Connection to Online Learning* (2013) by Abernethy, Chen, and Wortman-Vaughan.

Further bibliography for proper scoring rules:

- *Verification of Forecasts Expressed in Terms of Probability* (1950) by Glenn Brier.
- *Elicitation of personal probabilities and expectations* (1971) by Leonard Savage.
- *Strictly proper scoring rules, prediction, and estimation* (2007) by Gneiting and Raftery.