

Notes on Logarithmic Market Scoring Rules

Victor Shnayder Mike Ruberry

October 18, 2010

1 Overview

Note: This presentation will be relevant to the homework. We encourage you to pay attention, participate, and ask questions if anything is unclear.

Last week, we started talking about prediction markets, but didn't get into detail on how to implement them effectively. Today, we'll take a look at an approach called Market Scoring Rules, which was invented by Robin Hanson around 2002. We'll see that the main advantage of market scoring rules over a simple double auction based prediction market is that they incentivize people to participate.

But first, an aside: You may be wondering, who is Robin Hanson? Well, he's a professor at George Mason University, and he really believes that prediction markets are a powerful idea. In fact, he thinks that the country should be run as a futarchy—a form of government in which elected officials define measures of national welfare and prediction markets are used to determine which policies will have the most positive effect. We can discuss later whether we think this would be a good idea...

2 Slide 14 Introducing MSRs

Ok, so let's jump into what happens when we add "market" to "scoring rule".

3 Slide 15

Question: Form of these rules looks strange: a_i s depend on outcome. Is this a problem?

Answer: No—it doesn't depend on your report. Analogy to election, give everyone \$100 if Obama wins. Weird.

4 Slide 16

Question: Review: what is a proper scoring rule?

Answer: Setup: we want to obtain a probability distribution over some random variable, which will be publically realized in the future.

Scoring rule: A formula that determines an agent's reward as a function of their prediction of a probability distribution, compared with the the actual outcome. A proper scoring rule maximizes the expected reward when the agent reports truthfully.

5 Slide 17

Question: So, what's a *market* scoring rule?

Answer: We take a proper scoring rule, and allow people to make sequential reports against the same rule. To make a new report, an agent is required to pay the payoff of the previous report, once the final outcome is known.

Question: Does this still incentivize truthful reporting? Why?

Answer: Yes! Because the agent can't affect the payoff for the previous report, so they just want to maximize the expected payoff of their own report, just as before.

Question: Does this incentivize players to report in the first place?

Answer: It depends on the exact parameters, but if the original scoring rule is individually rational to participate in, then so is this: if the agent's prediction is different from the current report, they will expect to get a positive payoff from a truthful report of their prediction for the same reason that they would want to fix a mistaken report before it became official.

6 Slide 20 - Example

Imagine that we're interested in predicting the probability of the following: "In the year 2525, (what is the probability that) man is still alive? ". We will use the logarithmic scoring rule $x_i = \log(p_i)$, and initialize the probability at

0.5. Imagine that there is a sequence of agents with private beliefs 0.3 and 0.7. Let's look at their expected payoffs.

7 Slide 21

Question: if “man is still alive” happens, what are the payoffs?

Answer: Shown on slide22

8 Slide 22

Question: If we use the log scoring rule, do the payoffs have to be exactly as shown?

Answer: No—could use the a_i and b parameters

Question: As a joke, what is the only rational report to this prediction market?

Answer: 1. Why? If “man is not still alive”, it'll be difficult to make payoffs, so there is no case 2.

Question: Is the loss of the center bounded?

Answer: Yes. It corresponds to the maximum loss from any single report

Question: Is the loss/gain for a participant bounded?

Answer: Yes, for a single report, but no if agents are allowed to make an arbitrary number of reports. We'll see an example later.

9 Slide 24

Question: Why is it unreasonable for agents to have unbounded loss or gain in practice?

Answer: We only have a handwavy story: 1. They will run out of money. 2. They will update their beliefs over time and converge on something in the middle. Note that this is not explicitly modeled, but is implicitly assumed in any market.

10 Slide 25

Question: Why do we think that market scoring rules are better than simple information markets and using separate proper scoring rules for each agent?

Answer: Info markets: no trade theorem. Thin market- \rightarrow no trade- \rightarrow no info.

Separate scoring rules: works ok if asking one expert many questions, but if we get 100 answers to one question, it's unclear how to aggregate them (social choice, voting rules, etc.) We think that via the handwavy story above, people will converge on some reasonable aggregation of everyone's beliefs.

11 slide 26 – Victor's turn!

Question: So far, the discussion of market scoring rules has allowed any proper scoring rule to be used. Hanson argues that logarithmic scoring rules are particularly good. Why?

Answer: People has an easier time thinking about conditional probs. The log scoring rule preserves conditional probabilities.

Answer: Cheap-costs no more to elicit information about all combinations of base events than to elicit info just about the base events. Of course, this means that there is proportionally less incentive to provide that info.

Let's look at an example that involves conditional probabilities, and then look at a problem that motivates the market-like view of a MSR.

12 Slide 27

The French Laundry is a fancy restaurant in CA that is hard to get into.

Question: If you were to go, do you think you'd get a table? What matters? Can you give me a single probability that answers this question?

13 Slide 28

Answer: There are lots of relevant factors-getting a single number is tough

14 Slide 30

If we know whether or not you have a reservation, that makes the problem much simpler. We may be able to express the conditionals without having a good handle on the prior.

15 Slide 31

We like the log scoring rule because it is the only scoring rule that allows us to bet on conditionals without incurring extra cost, or making unexpected changes to the probabilities of events that we are not betting on. For example, if I bet on Table — Reservation, the log scoring rule ensures that $P(\text{reservation})$ doesn't change, which is as we would want.

Look: theorems

16 Slide 32

Question: Any potential problems here?

Answer: User interface, communication costs, cognitive costs (?), computation issues.

17 Slide 33 – Equivalent view as a market

Question: Why might we not like the model as presented above?

Answer: People may not find reporting distributions natural. Also, may want to directly express beliefs about just one part of the distribution. For example: “I think p_3 should be 20%. I don't have opinions about the others”. Of course, in the probability reporting setting, the agent could do the appropriate scaling to adjust appropriately, but it's not convenient. Also, people like money.

Question: What is the correspondence between reporting probability distributions and buying and selling shares?

18 Slide 34

Question: For the prediction problem from before, what are the contracts that will be traded in the market?

Answer: “\$1 if Man is alive in 2525”, etc

19 Futarchy?

So, do we like this? Is this useful in some contexts? Is it a good way to run the country?

20 Questions from the comments

- “I do not understand most of the terminology used in this paper. It feels like they are trying to present things intuitively, but by using terms that don’t seem to be defined.” **Yes!**
- Why would we need market scoring rules exactly? Prediction markets use current contract price and spread to drive accurate prediction, why use something else? People participate in them already... **Answer:** It is interesting that it can give a good probability estimate with very few traders (possibly only one).
- What if agents are risk averse? Risk taking?
- Collusion/manipulation?
- How do the different types of scoring rules (log, spherical, power, quadratic scoring rules) actually differ in application?
- I don’t really see the advantages that a market scoring rule offers over the peer prediction sequential reporting described in the Miller/Resnick paper, other than the fact that the market maker publicly reveals the current (latest) bet. What is the difference?
- Why is conditional independence so good? If other probs change, other agents can correct that through trading. Is it just the bounded loss for all events?

- Is the market model better for risk averse agents? Can buy fewer contracts—not change price as much, take on less risk. **Answer:** What about the equivalence? Can just report a closer prob in the other version—less profit, but less risk as well.
- How do you force agents to play in a Market prediction setting? Won't they be reluctant if they don't have any additional information (or edge)? **Answer:** Yes—if they have no extra info, they shouldn't participate
- A large b makes the market insensitive. Thus, it may not be possible for cash-constrained traders to move market sufficiently to the proper predictions. **Question:** Is this a bug or a feature?
- How does short-selling work in such a market?

Research idea: LMSR with center profit: charge per transaction? Other readings? —

However, as we shall see on Wednesday, these incentives really only apply to myopic traders who don't take into account how current actions may influence future price movements and thus profit opportunities.

Agents are loss averse. As a result, they will not make infinitesimally small adjustments in the predictions because their reference point will be the prevailing prediction. If they guess wrong, they will lose more utility than if they guessed right. Thus, it is possible that the mechanism will have a strong default bias, where the prediction will not adjust completely.

The way the blog phrased Hanson's results made me curious as to what would happen if we allow the market maker to participate in the market itself (I'm believe that this just means that the market maker can ensure 0 loss assuming he participates fairly).

I'm also interested in alternative rules for a market maker to follow. For instance, we could arrange such that the traders on the winning side split the profits evenly amongst themselves. How would such a payoff system affect the dynamics of the trading?

Also, in a continuous versions of LMSR, it seems as though the market maker would be exposed to unbounded potential loss unless he placed some sort of limit on the total number of shares that can be outstanding. Then the infinite sum we'd be taking is just over a finite support and is therefore bounded. In realistic terms this might amount to limiting liquidity, I think.

Finally, how does short-selling work in such a market? I can conceive of a contract based system that would make sense, but that seems to defeat the niceness of the market scoring rule with respect to doing away with trading contracts. In particular, it might work like this: Alice would tell Bob 'here, go ahead and pretend that you have these shares at the current market price [up to the definition of current being infinitesimally true, etc], and I will make good on actually giving them to you at some point in the future.' And then Alice hopes that at the future time the cost of the quantity of shares Bob was willing to take happens to be less than what Bob paid for them at the outset.