On Elicitation and Mechanism Design

Bo Waggoner, UPenn

January 2017

A hook

MONTHLY WEATHER REVIEW

EDITOR, JAMES E. CASKEY, JR.

Volume 78
Number 1

JANUARY 1950

Closed March 5, 1950
Issued April 15, 1950

VERIFICATION OF FORECASTS EXPRESSED IN TERMS OF PROBABILITY

GLENN W. BRIER

U. S. Weather Bureau, Washington, D. C.

[Manuscript received February 10, 1950]

INTRODUCTION

Verification of weather forecasts has been a controversial subject for more than a half century. There are a number of reasons why this problem has been so perplexing to meteorologists and others but one of the most important difficulties seems to be in reaching an agreement on the specification of a scale of goodness for weather forecasts.

numerically have been discussed previously [1, 2, 3, 4] so that the purpose here will not be to emphasize the enhanced usefulness of such forecasts but rather to point out how some aspects of the verification problem are simplified or solved.

VERIFICATION FORMULA

Suppose that on each of n occasions an event can occur

Goals of this talk

1. Advertise **information elicitation**
2. Point out (deep) connections to **mechanism design**
3. (Optional) discuss some research

Information elicitation - overview

1. Agent makes report r
2. Mechanism observes outcome ω of some event
3. Mechanism assigns score $S(r, \omega)$

Assumption: Agent with belief q reports

$$\arg\max_r \mathop{\mathbb{E}}_{\omega \sim q} S(r, \omega).$$

Question: How does the report depend on S and q?

In other words: how to design S to elicit certain properties?

Examples for scalar outcomes and reports

Event ω in $\mathbb{R}$, e.g. inches of rain tomorrow.
Report r in $\mathbb{R}$.

- $S(r, \omega) = -(r - \omega)^2$. *mean*
- $S(r, \omega) = -|r - \omega|$. *median*
- $S(r, \omega) = \mathbf{1}[r = \omega]$. *mode – if finite outcome space*
- $S(r, \omega) = e^r\,(1 + \omega - r)$. *mean*

Example impossibility result

Fact: There is no S such that the agent reports the **variance**.

$$Var(q) = \arg\max_r \mathop{\mathbb{E}}_{\omega \sim q} S(r, \omega).$$

Lemma

If report r is optimal for both beliefs p and q, then it is also optimal for belief $\alpha p + (1-\alpha)q$.

Proof of fact: Consider distributions on $\{0, 1\}$ ....

Solutions: (1) elicit multidimensional response;
(2) draw multiple samples (work in progress).

Proper scoring rules: report $\in \Delta_\Omega$

Fact (McCarthy 1956; Savage 1971; Gneiting and Raftery 2007)

(1) For every convex $g : \Delta_\Omega \to \mathbb{R}$*, there exists an (easy-to-construct) proper scoring rule* S_g *with* $\mathbb{E}_{\omega \sim q} S_g(q, \omega) = g(q)$.

(2) All proper S *are of the above form.*

Key idea:

1. Expected score is **linear** in "type" q
2. Agent selects max of these linear functions
   $\implies$ convexity of g

Truthful interim allocation rules

In mechanism design, fix reports of all others and consider i's problem.

1. Agent reports values $\hat{v}_a$ per allocation a
2. Mechanism outputs distribution $M(\hat{\mathbf{v}})$ over $\mathcal{A}$
3. Agent utility is $\mathbb{E}_{a \sim M(\hat{\mathbf{v}})} v_a$

Key idea:

1. Expected utility is **linear** in "type" v
2. Agent selects max of these linear functions
   $\implies$ convexity of utility as function of type

My point (if any)

IE and MD **share mathematical foundations**.[1]
Utilities are convex in type; allocations are subgradients.

- e.g. characterizations: weak monotonicity $\leftrightarrow$ power diagrams
  Saks and Yu 2005 $\leftrightarrow$ *Lambert et al. 2008*

- e.g: the rest of the talk

[1] Key paper: Frongillo and Kash 2014.

Some questions in foundations

Room to explore...

- Role of convexity in truthful mechanism design (Zihe's talk)
- "Restricted" or "compressed" truthfulness
  reporting only part (a "property") of a large type
- Work on multidimensional setting together
  seems very hard in both areas... for connected reasons
  - In 1d: view Lambert et al. 2008 as extension of Myerson's lemma

An ad auction problem

Based on *Designing Mechanisms for Daily Deals* by Yang Cai, Mohammad Mahdian, Aranyak Mehta, and Bo Waggoner (WINE 2013).

- Single ad-auction slot to allocate
- Each bidder has **value** v_i and **click-probability** p_i (both private information)
- Website user's utility for seeing ad with p is $g(p)$

Goal: Maximize social welfare: $v_{i^*} + g(p_{i^*})$.
Tool: Can pay bonuses if clicked.

Suppose we knew click-through rates

1. Pick i^* maximizing welfare
2. Pay her $g(p_{i^*})$
3. Charge her the second-highest $v_i + g(p_i)$

Utility $= v_{i^*} + g(p_{i^*}) - [v_i + g(p_i)]$.

Fact (scoring rule characterization – recall)

For all convex g, there exists $S_g(p, \textit{click})$ where:

- *truthful reporting of p_{i^*} yields $g(p_{i^*})$ in expectation;*
- *any false report yields less.*

A scoring-rule solution

1. Pick i^* maximizing welfare
2. Pay her $S_g(p_{i^*}, \text{click})$
3. Charge her second-highest $v_i + g(p_i)$

$$\begin{aligned}\text{winner's utility} &= v_{i^*} + \mathbb{E}\, S_g(p_{i^*}, \text{click}) \;-\; [v_i + g(p_i)] \\ &\le v_{i^*} + g(p_{i^*}) \;-\; [v_i + g(p_i)]\end{aligned}$$

Truthfulness:

- Given that i^* wins, optimal to report p_{i^*} truthfully.
- i^* wants to win $\iff$ she maximizes the objective.

A VCG generalization

- Mechanism has a set of allocations A
- For each allocation a, agent i has:
  - value $v_i(a)$
  - belief $p_{i,a}$ over some space $\Omega_{i,a}$

Theorem

There is a mechanism to maximize

$$\sum_i v_i(a) \; + \; g_a(p_{1,a}, \ldots, p_{n,a})$$

if and only if each g_a is **component-wise convex**.

Payment rule for i: VCG payment minus $S_{g_{a^*},-i}(p_i, \omega_{i,a^*})$.

Example application (1/2)

Idea: purchase a route from s to t in a network.
Each edge e is controlled by an agent with a

- cost c_e for edge utilization
- distribution p_e over travel time along the edge

Objective: pick a path $\mathbf{a}$ maximizing

$$g\left(\mathbf{p}_{\mathbf{a}}\right) \; - \; \sum_{e \in \mathbf{a}} c_e$$

where

- $\mathbf{p}_{\mathbf{a}}$ is the distribution on total travel time
- g is convex (modeling **risk aversion**)

Example application (2/2)

Idea: allocate bandwidth to communicators.
Each i has value v_i for sending a message X_i.

$X_1, \dots, X_n$ all drawn jointly from p known to all agents (not designer).

Objective: pick a subset S maximizing

$$\sum_{i \in S} v_i \;-\; H(X_i : i \in S).$$

Not well-defined: what if agents disagree on p?

Future directions

- Concretely: practicality problems with this mechanism
- More problems where type = (value, belief)
- Mechansims as **aggregations** of preferences and beliefs.

…and a whole world of info. elicitation problems!

Thanks!