Useful Forecasting

Belief Elicitation for Decision Making

Lars Wittrock

May 2, 2022

Maastricht University

Introduction

Introduction

Example: Obama and the hunt for Bin Laden in 2011
(well documented by Friedman and Zeckhauser 2014)

- Decision to attack or wait has to be made by Obama.
- Obama has no knowledge himself if Bin Laden is at the suspected location.
- All agents are asked to report a probability estimate of Bin Laden being at the suspected location.
- Many agents strategically misreported their belief to influence Obama's final decision.

Introduction

Further examples:

- Politicians deciding on Covid-19 measures based on advice from a group of experts.
- A manager in a firm deciding on which project to pursue based on information from the respective technical/sales departments.
- ...

Open question: How can a principal best elicit beliefs?

Introduction

- Mechanisms based on a (proper) scoring rule are the main tool to elicit beliefs.
    - Single person: QSR, BSR, ...
    - Group: Prediction markets and prediction polls
- Scoring rule mechanisms make unrealistic assumptions:
    - The elicited belief is not used to make a decision, or
    - Experts care only about the (monetary) payoff from the mechanism.

Research Questions

- How can a principal incentivize experts to report their belief truthfully?
- What is the best mechanism if the principal can only consult a single expert?

Literature

- Scoring rules and mechanism design (Gneiting and Raftery 2007 and Conitzer 2009)
- Elicited beliefs are used for decision making. Experts are decision-agnostic. (Berg and Rietz 2003, Oesterheld and Conitzer 2019, Othman and Sandholm 2010, Chen and Kash 2011, Chen, Kash, et al. 2011 and Dimitrov and Sami 2010)
- Elicited beliefs are used for decision making and experts have decision preferences. The principal has knowledge of the expert's action preferences. (Boutilier 2012)
- Decentralized decision making and strategic information transmission. (Holmström 1977, Holmström 1984 and Crawford and Sobel 1982)

Model

Model

Situation

- Principal is faced with a choice between two actions: $A = \{a_1, a_2\}$.
- Two states of the world: $\Omega = \{\omega_1, \omega_2\}$.
- Principal has state-dependent preferences over the actions.
- The state is revealed after the action choice.

Experts:

- n different risk-neutral experts.
- Each expert has a private belief, $\mu_i \in [0, 1]$, about the state being ω_2.
- Each expert has unobservable action preferences: $U_i(a_2) := u_i$.
- Each expert knows the principal's preferences.

Model

Principal:

- No information about the state of the world.
- Ask each expert to report a belief, r_i.
- Principal forms a belief equal to the mean of all reported beliefs, $\mu^P := \bar{r} = \frac{1}{n}\sum_{i \in N} r_i$.
- Choose:

$$\mathcal{D} = \begin{cases} a_2 & \text{if } \bar{r} \geq \alpha \\ a_1 & \text{if } \bar{r} < \alpha \end{cases}$$

Question:

- How can the principal elicit truthful reports from the expert(s)?

Background on Scoring Rules

A scoring rule is a function $S : [0,1] \times \Omega \to \mathbb{R}$ which determines a monetary payoff $S(r, \omega)$ based on the reported belief $r \in [0,1]$ and the state of the world ω.

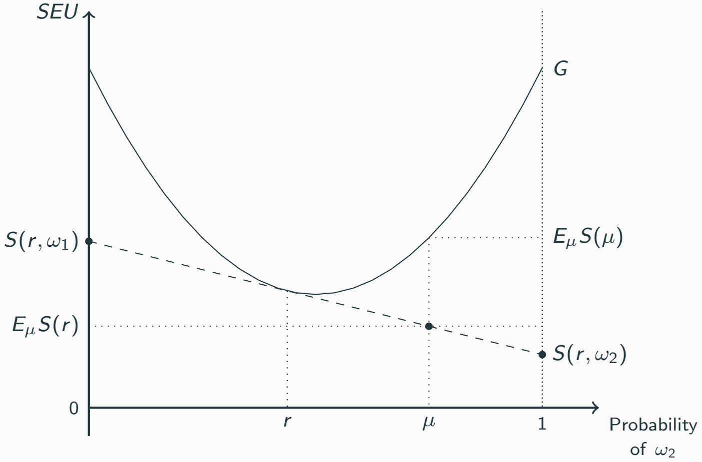

Single Expert

Expert Behavior

- Subjective Expected Utility:

$$SEU(r) = \begin{cases} E_\mu S(r) + u & \text{if } r \geq \alpha \\ E_\mu S(r) & \text{if } r < \alpha \end{cases}$$

- Trade-off between benefit and cost of misreporting if true belief would lead to less preferred action.
- Optimal report in terms of true belief (with $u > 0$):

$$r^* = \begin{cases} \mu & \text{if } \mu \notin [c_-, \alpha] \\ \alpha & \text{if } \mu \in [c_-, \alpha] \end{cases}$$

- c_-: $u = E_{c_-} S(c_-) - E_{c_-} S(\alpha)$.

Expert Behavior

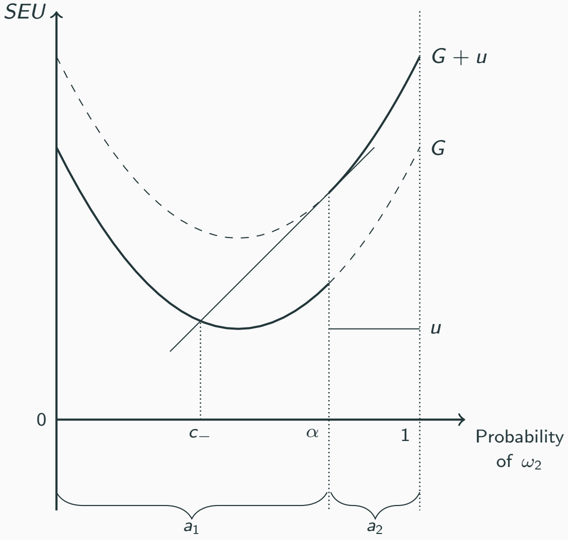

Theoretically Optimal Mechanism

Theorem

For any belief, μ, and some fixed outside preferences, u, truth-telling is a dominant strategy if and only if the scoring rule is given by S^* with

$$S^*(r,\omega) = \begin{cases} S(r,\omega) & \text{if } r \geq \alpha \\ S(r,\omega) + u & \text{if } r < \alpha \end{cases}$$

where $S(r,\omega)$ is any proper scoring rule.

Theoretically Optimal Mechanism

Theorem

For any belief, μ, and some fixed outside preferences, u, truth-telling is a dominant strategy if and only if the scoring rule is given by S^* with

$$S^*(r,\omega) = \begin{cases} S(r,\omega) & \text{if } r \geq \alpha \\ S(r,\omega) + u & \text{if } r < \alpha \end{cases}$$

where $S(r,\omega)$ is any proper scoring rule.

- Truth-telling is guaranteed if a mechanism takes into account the expert's action preferences.
- No mechanism exists that can guarantee truthful reporting if action preferences are unobservable.

Best Practical Mechanism

Definition

Best practical mechanism:

- It is feasible, i.e. $S(r, \omega) \in [0, B] \; \forall \; r, \omega$ and
- It minimizes the set of types that would misreport.

Best Practical Mechanism

Definition

Best practical mechanism:

- It is feasible, i.e. $S(r, \omega) \in [0, B] \ \forall \ r, \omega$ and
- It minimizes the set of types that would misreport.

Theorem

The best practical mechanism is given by the scoring rule S^* such that:

$$S^*(r, \omega_1) = \begin{cases} B & \text{if } r < \alpha \\ 0 & \text{if } r \geq \alpha \end{cases}$$

and

$$S^*(r, \omega_2) = \begin{cases} 0 & \text{if } r < \alpha \\ 2B(1 - \alpha) & \text{if } r \geq \alpha \end{cases}$$

Best Practical Mechanism

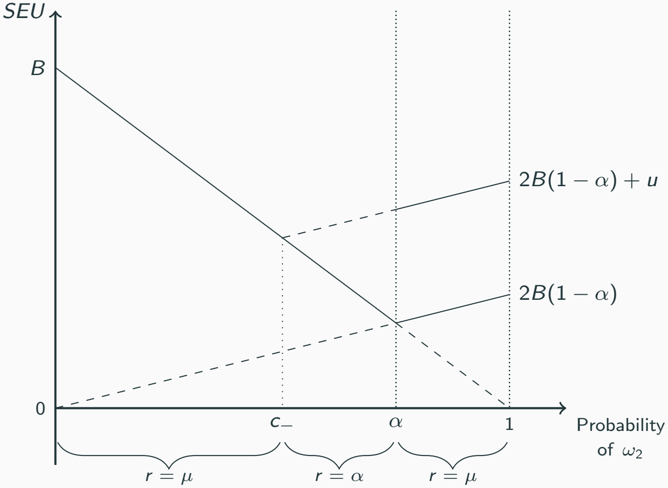

Multiple Experts

Direct Reporting Mechanism

- All experts independently report a belief to the principal.
- Beliefs are aggregated by simple mean: $\bar{r} = \frac{\sum_{i \in N} r_i}{n}$
- The principal announces the following decision rule:

$$\mathcal{D} = \begin{cases} a_2 & \text{if } \bar{r} \geq \alpha \\ a_1 & \text{if } \bar{r} < \alpha \end{cases}$$

- Subjective expected utility of each expert:

$$SEU_i(r_i) = \begin{cases} E_{\mu_i} S(r_i) + u_i & \text{if } \frac{1}{n} r_i + \frac{n-1}{n} \tilde{r}_{-i} \geq \alpha \\ E_{\mu_i} S(r_i) & \text{if } \frac{1}{n} r_i + \frac{n-1}{n} \tilde{r}_{-i} < \alpha \end{cases}$$

Expert Behavior

Definition

Expert i is considered to be pivotal if $\tilde{r}_{-i}$ is such that $\frac{n-1}{n}\tilde{r}_{-i} < \alpha \leq \frac{n-1}{n}\tilde{r}_{-i} + \frac{1}{n}$.

Observation 1

Given some $\tilde{r}_{-i}$, if expert i is not pivotal, for any (strictly) proper scoring rule S it is (strictly) optimal for the expert to report his belief truthfully, $r_i = \mu_i$.

Expert Behavior

Defining the pivotal report, such that $\bar{r} = \alpha$:

$$c_{i,+} := \alpha + (n-1)(\alpha - \tilde{r}_{-i})$$

Observation 2

Given some $\tilde{r}_{-i}$, μ_i and u_i, the only report that can be optimal is $r_i = \mu_i$ or $r_i = c_{i,+}$.

Expert Behavior

Defining the pivotal report, such that $\bar{r} = \alpha$:

$$c_{i,+} := \alpha + (n-1)(\alpha - \tilde{r}_{-i})$$

Observation 2

Given some $\tilde{r}_{-i}$, μ_i and u_i, the only report that can be optimal is $r_i = \mu_i$ or $r_i = c_{i,+}$.

Optimal report (with $u > 0$) is given by:

$$r_i^* = \begin{cases} \mu_i & \text{if } c_{i,+} \notin [0,1] \\ \begin{cases} \mu_i & \text{if } \mu_i \notin (c_{i,-}, c_{i,+}] \\ c_{i,+} & \text{if } \mu_i \in (c_{i,-}, c_{i,+}] \end{cases} & \text{if } c_{i,+} \in [0,1] \end{cases}$$

Main Result

Theorem

For any number of experts ($n \geq 2$), any strictly proper scoring rule S, all experts reporting their belief truthfully, $r_i = \mu_i \ \forall \ i$, is the unique and strict Nash equilibrium if,

1) **Diversity:** *the profile of action preferences is not such that $\forall \ i \ \ u_i \geq 0$ or vice versa, and*

2) **No pivotality:** $\bar{\mu} \notin [\alpha - \frac{1}{n}, \alpha + \frac{1}{n})$.

Discussion and Summary

Alternative Mechanisms

Other mechanisms:

- Sequential reporting
- Prediction markets
- Simple voting

Other methods of aggregating beliefs:

- Median beliefs
- Some weighted average

Summary

Results:

- No mechanism exists that makes truthful reporting a dominant strategy.
- In the single expert setting it is best to delegate the decision to the expert.
- With multiple experts, truth-telling is the unique Nash Equilibrium under two conditions: Preference diversity and no pivotality.

Summary

Results:

- No mechanism exists that makes truthful reporting a dominant strategy.
- In the single expert setting it is best to delegate the decision to the expert.
- With multiple experts, truth-telling is the unique Nash Equilibrium under two conditions: Preference diversity and no pivotality.

Open questions:

- True state only revealed after a certain action choice?
- Correlated beliefs and/or preferences?
- 3 or more states/actions

Questions?

References

Berg, Joyce E. and Thomas A. Rietz (2003). "Prediction Markets as Decision Support Systems". In: *Information Systems Frontiers* 5.1, pp. 79–93. DOI: 10.1023/a:1022002107255.

Boutilier, Craig (2012). "Eliciting Forecasts from Self-Interested Experts: Scoring Rules for Decision Makers". In: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*. AAMAS '12. Valencia, Spain: International Foundation for Autonomous Agents and Multiagent Systems, pp. 737–744. ISBN: 0981738125.

Chen, Yiling and Ian Kash (2011). "Information Elicitation for Decision Making". In: *AAMAS '11: The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*.

Chen, Yiling, Ian Kash, et al. (Dec. 2011). "Decision Markets With Good Incentives". In: *WINE*, pp. 72–83. DOI: 10.1007/978-3-642-25510-6_7.

Conitzer, Vincent (2009). "Prediction Markets, Mechanism Design, and Cooperative Game Theory". In: *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*. UAI '09. Montreal, Quebec, Canada: AUAI Press, pp. 101–108. ISBN: 9780974903958.

Crawford, Vincent P. and Joel Sobel (Nov. 1982). "Strategic Information Transmission". In: *Econometrica* 50.6, p. 1431. DOI: 10.2307/1913390.

Dimitrov, Stanko and Rahul Sami (2010). "Composition of markets with conflicting incentives". In: *EC '10*.

Friedman, Jeffrey and Richard Zeckhauser (Dec. 2014). "Handling and Mishandling Estimative Probability: Likelihood, Confidence, and the Search for Bin Laden". In: *Intelligence and National Security* 30, pp. 77–99. DOI: 10.1080/02684527.2014.885202.

Gneiting, Tilmann and Adrian E. Raftery (2007). "Strictly Proper Scoring Rules, Prediction, and Estimation". In: *American Statistical Association*.

Holmström, B. (1977). "On incentives and control in organizations". In.

— (1984). "On the Theory of Delegation". In: ed. by M. Boyer and R. Kihlstrom. North Holland. Chap. Bayesian Models in Economic Theory, pp. 115–141.

Oesterheld, Caspar and Vincent Conitzer (2019). "Eliciting information for decision making from individual and multiple experts". In.

Othman, Abraham and Tuomas Sandholm (Jan. 2010). "Decision Rules and Decision Markets". In: vol. 1, pp. 625–632.