Non-distortionary belief elicitation

Marcin Pęski Colin Stewart

December 9, 2024

Introduction

Experiments often include belief elicitation.

- examples: testing belief-dependent models, cognitive uncertainty,
- incentives improve accuracy (Schlag et al. (2015) and many others),
- range of methods to incentivize, including binarized Becker-DeGroot-Marschak (BDM) scoring rule (Hossain and Okui (2013)),
- incentivizing scoring rule does not need to be explicit (Danz et al. (2022)).

Cognitive load constraints make elicitation of "all" beliefs difficult.

Often, the researcher is interested in *action-dependent* questions:

- what is your expected payoff?
- how likely would you like to change your action after learning the state of the world?
- how much would you pay for the option to change your action?

Example

Subjects write a test in a classroom. The test has two parts (Micro, and Macroeconomics) with 20 multiple choice questions each.

At the end of the test, each subject is asked:

- How many questions did you answer correctly?
- How many more correct questions do you have in Macro vs Micro part?
- How likely are you to cross 50% threshold?

- Questions are action-dependent.
- In practice, difficult to elicit the same information using action-independent questions.

Introduction

Incentivizing elicitation of action-dependent beliefs may distort behavior in the original problem:

- "contamination" in (Healy, 2024),
- "moral hazard" (Chambers and Lambert, 2021),
- sometimes, elicitation is carefully to avoid distortion (Hu, 2023), (Gaechter and Renner, 2010)

> *The elicitation of beliefs and utilities may alter game play. And playing the game may alter the beliefs and utilities that subjects report. ... We are not aware of any way to remove contamination through the experimental design, so instead we embrace it. (Healy, 2024).*

Research goals

What action-dependent question can be incentivized without distorting incentives in the original experiment?

Answer

Expected payoffs or "some" affine transformations thereof.

Research goals

What action-dependent question can be incentivized without distorting incentives in the original experiment?

Answer

Expected payoffs or "some" affine transformations thereof.

Outline

Model

Decision problem

$$\max_a \sum_{\theta} p(\theta) u(a, \theta)$$

- Single DM,
- Actions $a \in A$,
- States $\theta \in \Theta$, beliefs $p \in \Delta\Theta$
- Payoffs can be binarized, i.e, $u : A \times \Theta \longrightarrow [0, 1]$.
- We assume that none of the actions are dominated.

Model

Question

- Action-dependent question: $X(a, \theta) \in \mathbb{R}$.
- Researcher asks about the expected value $r = \mathbb{E}_p X(a, .)$,
  - "linear" belief.

- Action-dependent question: $X(a,\theta) \in \mathbb{R}$.
- Researcher asks about the expected value $r = \mathbb{E}_p X(a,.)$.

Example

1. Expected payoffs: $X(a,\theta) = u(a,\theta)$
2. Expected regret: $X(a,\theta) = u(a,\theta) - \max_b u(b,\theta)$
3. (*Ex post*) correct choice: $X(a,\theta) = \begin{cases} 1 & \text{if } a \in \arg\max_{b\in A} u(b;\theta) \\ 0 & \text{otherwise.} \end{cases}$
4. Probability of state θ_0: $X(a,\theta) = 1\{\theta = \theta_0\}$

- Single (1-dimensional) question:
  - cognitive load,
  - with $|\Theta| - 1$ questions, everything is incentivizable.
- *Linear* property of beliefs, $\mathbb{E}_p X(a;.)$: practical interest, but restrictive
- Incentvization through scoring rule:

$$\max_{a,r} V(a,r,\theta),$$

where payoffs may combine original payoffs and belief-incentivizing scoring rule V_0,

$$V(a,r,\theta) = (1-\alpha)u(a,\theta) + \alpha V_0(a,r,\theta).$$

Model

Question

- Single (1-dimensional) question:
  - cognitive load,
  - with $|\Theta| - 1$ questions, everything is incentivizable.
- *Linear* property of beliefs, $\mathbb{E}_p X(a;.)$: practical interest, but restrictive
- Incentvization through scoring rule:

$$\max_{a,r} V(a, r, \theta),$$

where payoffs may combine original payoffs and belief-incentivizing scoring rule V_0,

$$V(a, r, \theta) = (1 - \alpha)u(a, \theta) + \alpha V_0(a, r, \theta).$$

- Single (1-dimensional) question:
  - cognitive load,
  - with $|\Theta| - 1$ questions, everything is incentivizable.
- *Linear* property of beliefs, $\mathbb{E}_p X(a;.)$: practical interest, but restrictive
- Incentvization through scoring rule:

$$\max_{a,r} V(a, r, \theta),$$

where payoffs may combine original payoffs and belief-incentivizing scoring rule V_0,

$$V(a, r, \theta) = (1-\alpha)u(a, \theta) + \alpha V_0(a, r, \theta).$$

Incentivizability

Question X is *incentivizable* if there exists a scoring rule V such that

$$\arg\max_{a,r} \mathbb{E}_p V(a, r, .) = \left\{ (a, \mathbb{E}_p X(a; .)) : a \in \arg\max_{b \in A} \mathbb{E}_p u(b; \cdot) \right\}.$$

- strict incentives for reporting $\mathbb{E}_p X(a; .)$,
- without distorting the behavior in original problem.

Incentivizability

Question X is *incentivizable* if there exists a scoring rule V such that

$$\arg\max_{a,r} \mathbb{E}_p V(a, r, .) = \left\{ (a, \mathbb{E}_p X(a; .)) : a \in \arg\max_{b \in A} \mathbb{E}_p u(b; \cdot) \right\}.$$

- strict incentives for reporting $\mathbb{E}_p X(a; .)$,
- without distorting the behavior in original problem.

Outline

Sufficient conditions

Lemma

$X(a,\theta) = d(\theta)$ for any $d \in \mathbb{R}$ is incentivizable.
$X(a,\theta) = u(a,\theta) + d(\theta)$ for any $d \in \mathbb{R}$ is incentivizable.

Questions about payoffs plus an action-independent variable can be incentized.

Proof.

BDM: Take any $x_{\min} < u(a;\theta) + d(\theta) < x_{\max}$ and let

$$V(r,a,\theta) = \frac{1}{x_{\max} - x_{\min}} \int_{x_{\min}}^{r} X(a;\theta)\,dx + \frac{1}{x_{\max} - x_{\min}} \int_{r}^{x_{\max}} x\,dx$$

$$= \frac{1}{x_{\max} - x_{\min}} \left[\left((u(a;\theta) + d(\theta))(r - L) - \frac{1}{2} r^2 \right) + \frac{x_{\max}^2}{2}, \right]$$

Sufficient conditions

Lemma

$X(a,\theta) = d(\theta)$ for any $d \in \mathbb{R}$ is incentivizable.
$X(a,\theta) = u(a,\theta) + d(\theta)$ for any $d \in \mathbb{R}$ is incentivizable.

Questions about payoffs plus an action-independent variable can be incentized.

Proof.

BDM: Take any $x_{\min} < u(a;\theta) + d(\theta) < x_{\max}$ and let

$$V(r,a,\theta) = \frac{1}{x_{\max} - x_{\min}} \int_{x_{\min}}^{r} X(a;\theta)\,dx + \frac{1}{x_{\max} - x_{\min}} \int_{r}^{x_{\max}} x\,dx$$

$$= \frac{1}{x_{\max} - x_{\min}} \left[\left((u(a;\theta) + d(\theta))(r - L) - \frac{1}{2} r^2 \right) + \frac{x_{\max}^2}{2}, \right]$$

□

Sufficient conditions

Lemma

For any question X, any $\gamma, \kappa : A \longrightarrow \mathbb{R}$, let $Y(a, \theta) = \gamma(a)X(a, \theta) + \kappa(a)$. If X is incentivizable, then Y is incentivizable.

Affine transformations of incentivizable questions can be incentivized.

Proof.

Take $V_Y(a, r, \theta) = V_X(a, \frac{1}{\gamma(a)}(r - \kappa(a)), \theta)$. □

Sufficient conditions

Lemma

For any question X, any $\gamma, \kappa : A \longrightarrow \mathbb{R}$, let $Y(a, \theta) = \gamma(a)X(a, \theta) + \kappa(a)$. If X is incentivizable, then Y is incentivizable.

Affine transformations of incentivizable questions can be incentivized.

Proof.

Take $V_Y(a, r, \theta) = V_X(a, \frac{1}{\gamma(a)}(r - \kappa(a)), \theta)$. □

Sufficient conditions

Aligned representation

Question X is *aligned* with u on $B \subseteq A$ if and only if there are $\gamma, \kappa : B \longrightarrow \mathbb{R}$, and $d \in \mathbb{R}^{\Theta}$ such that for each $a \in B$

$$\begin{aligned} X(a,\theta) &= \gamma(a)\left(u(a,\theta) + d(\theta)\right) + \kappa(a), \textit{ or} \\ &= \gamma(a)\left(d(\theta)\right) + \kappa(a) \end{aligned}$$

- In the second case, we say that X is *trivial*.
- Corollary: Any X that is aligned on A is incentivizable.

$$X(a,\theta) = \gamma(a)\left(u(a,\theta) + d(\theta)\right) + \kappa(a), \text{ or}$$
$$= \gamma(a) d(\theta) + \kappa(a)$$

Example

1. ✓ Expected payoffs: $X(a,\theta) = u(a,\theta)$
2. ✓ Expected regret: $X(a,\theta) = u(a,\theta) - \max_b u(b,\theta)$
3. ✗ (*Ex post*) correct choice: $X(a,\theta) = \begin{cases} 1 & \text{if } a \in \arg\max_{b\in A} u(b;\theta) \\ 0 & \text{otherwise.} \end{cases}$
4. ✓ Probability of state θ_0: $X(a,\theta) = 1\{\theta = \theta_0\}$

Sufficient conditions
Notation

- Let $\bar{\mathsf{X}}(a,\theta) = \mathsf{X}(a,\theta) - \frac{1}{|\Theta|}\sum_{\theta'} \mathsf{X}(a,\theta')$.
- Let $\Delta_a^b(\theta) = \bar{\mathsf{u}}(b,\theta) - \bar{\mathsf{u}}(a,\theta)$.
- If X is aligned on B, then for all $a, b \in B$, there is x and y such that

$$\bar{\mathsf{X}}(a) = x\bar{\mathsf{X}}(b) + y\Delta_a^b$$

    - equivalence, if $|B| = 2$,
- If X is trivial, all $\bar{\mathsf{X}}(a)$ are collinear.

Outline

Necessary conditions

If a and b are best responses at the same belief, and there is no other optimal action, we say that a, b are *adjacent*.

Adjacency Lemma

If X is incentivizable, then X is aligned with u on each pair of adjacent actions $\{a, b\}$.

- Suppose X is incentivizable, a, b are adjacent, and p is a belief st. $\mathbb{E}_p u(a,.) = \mathbb{E}_p u(b,.)$,.
- If $r = \mathbb{E}_p X(a,.)$ and $s = \mathbb{E}_p X(b,.)$, then

$$\mathbb{E}_p V(a, r, .) = \mathbb{E}_p V(b, s, .).$$

- **Observation** The expected maximum payoff from scoring rule is affine over a set of beliefs if and only if the same choice is a best response.

Necessary conditions
Scoring rule

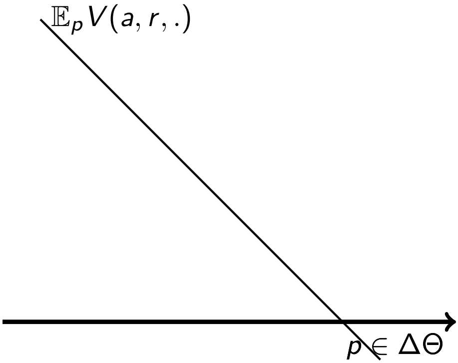

$\mathbb{E}_p V(a, r, .)$
$p \in \Delta\Theta$

Necessary conditions

Scoring rule

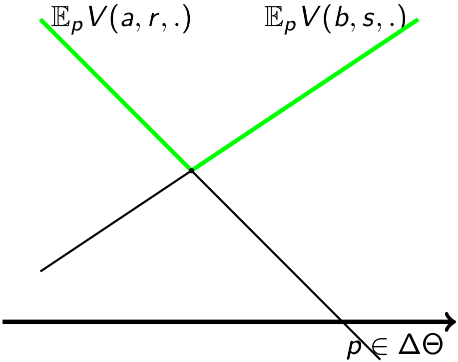

Necessary conditions

Scoring rule

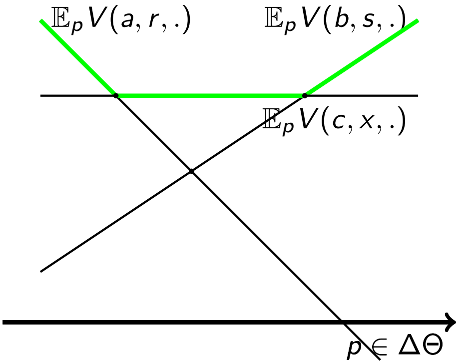

An action is the best response for a set of beliefs if and only if the expected maximum scoring payoff is affine over this set.

Necessary conditions

Scoring rule

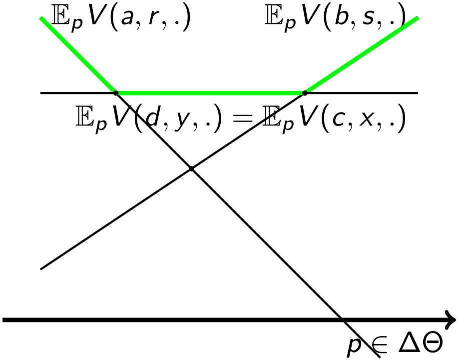

If another action is indifferent at one point in the interior of the set, it must be indifferent over the rest of the set.

- Hence, for all $p \in P_0 = \{q : a \in A(q), r = \mathbb{E}_q X(a, .)\}$
  - (a, r) is optimal at P_0
  - the expected maximal scoring payoff is equal to the expected payoff from (a, r), hence affine over P_0,
  - take $q \in \text{int} P_0$ and let $s = \mathbb{E}_q X(b, .)$,
  - Observation $\Rightarrow (b, s)$ is optimal at all $p \int P_0$,
  - $\mathbb{E}_p X(b, .) = s$ for all $p \int P_0$.
- $dp \perp 1, \mathrm{u}(b) - \mathrm{u}(a), \mathrm{X}(a) \Rightarrow dp \perp \mathrm{X}(b)$,
- Linear algebra
- $\Rightarrow \bar{\mathrm{X}}(b) \in \text{span}(\bar{\mathrm{X}}(a), \bar{\mathrm{u}}(b) - \bar{\mathrm{u}}(a))$,
- $=> X$ is aligned on $\{a, b\}$.

- Hence, for all $p \in P_0 = \{q : a \in A(q), r = \mathbb{E}_q X(a,.)\}$
  - (a, r) is optimal at P_0
  - the expected maximal scoring payoff is equal to the expected payoff from (a, r), hence affine over P_0,
  - take $q \in \text{int} P_0$ and let $s = \mathbb{E}_q X(b,.)$,
  - Observation $\Rightarrow (b, s)$ is optimal at all $p \int P_0$,
  - $\mathbb{E}_p X(b,.) = s$ for all $p \int P_0$.
- $dp \perp 1, \mathrm{u}(b) - \mathrm{u}(a), \mathrm{X}(a) \Rightarrow dp \perp \mathrm{X}(b)$,
- Linear algebra
- $\Rightarrow \bar{\mathrm{X}}(b) \in \text{span}(\bar{\mathrm{X}}(a), \bar{\mathrm{u}}(b) - \bar{\mathrm{u}}(a))$,
- => X is aligned on $\{a, b\}$.

- Hence, for all $p \in P_0 = \{q : a \in A(q), r = \mathbb{E}_q X(a,.)\}$
  - (a, r) is optimal at P_0
  - the expected maximal scoring payoff is equal to the expected payoff from (a, r), hence affine over P_0,
  - take $q \in \text{int} P_0$ and let $s = \mathbb{E}_q X(b,.)$,
  - Observation $\Rightarrow (b, s)$ is optimal at all $p \int P_0$,
  - $\mathbb{E}_p X(b,.) = s$ for all $p \int P_0$.
- $dp \perp 1, \mathrm{u}(b) - \mathrm{u}(a), \mathrm{X}(a) \Rightarrow dp \perp \mathrm{X}(b)$,
- Linear algebra
- $\Rightarrow \bar{\mathrm{X}}(b) \in \mathrm{span}(\bar{\mathrm{X}}(a), \bar{\mathrm{u}}(b) - \bar{\mathrm{u}}(a))$,
- $=> X$ is aligned on $\{a, b\}$.

- Hence, for all $p \in P_0 = \{q : a \in A(q), r = \mathbb{E}_q X(a, .)\}$
  - (a, r) is optimal at P_0
  - the expected maximal scoring payoff is equal to the expected payoff from (a, r), hence affine over P_0,
  - take $q \in \mathrm{int} P_0$ and let $s = \mathbb{E}_q X(b, .)$,
  - Observation $\Rightarrow$ (b, s) is optimal at all $p \int P_0$,
  - $\mathbb{E}_p X(b, .) = s$ for all $p \int P_0$.
- $dp \perp 1, \mathrm{u}(b) - \mathrm{u}(a), \mathrm{X}(a) \Rightarrow dp \perp \mathrm{X}(b)$,
- Linear algebra
- $\Rightarrow \bar{\mathrm{X}}(b) \in \mathrm{span}(\bar{\mathrm{X}}(a), \bar{\mathrm{u}}(b) - \bar{\mathrm{u}}(a))$,
- => X is aligned on $\{a, b\}$.

Necessary conditions

Adjacency Lemma

- Hence, for all $p \in P_0 = \{q : a \in A(q), r = \mathbb{E}_q X(a, .)\}$
  - (a, r) is optimal at P_0
  - the expected maximal scoring payoff is equal to the expected payoff from (a, r), hence affine over P_0,
  - take $q \in \mathrm{int} P_0$ and let $s = \mathbb{E}_q X(b, .)$,
  - Observation $\Rightarrow (b, s)$ is optimal at all $p \int P_0$,
  - $\mathbb{E}_p X(b, .) = s$ for all $p \int P_0$.
- $dp \perp 1, \mathrm{u}(b) - \mathrm{u}(a), \mathrm{X}(a) \Rightarrow dp \perp \mathrm{X}(b)$,
- Linear algebra
- $\Rightarrow \bar{\mathrm{X}}(b) \in \mathrm{span}(\bar{\mathrm{X}}(a), \bar{\mathrm{u}}(b) - \bar{\mathrm{u}}(a))$,
- => X is aligned on $\{a, b\}$.

- Hence, for all $p \in P_0 = \{q : a \in A(q), r = \mathbb{E}_q X(a,.)\}$
  - (a, r) is optimal at P_0
  - the expected maximal scoring payoff is equal to the expected payoff from (a, r), hence affine over P_0,
  - take $q \in \text{int} P_0$ and let $s = \mathbb{E}_q X(b,.)$,
  - Observation $\Rightarrow (b, s)$ is optimal at all $p \int P_0$,
  - $\mathbb{E}_p X(b,.) = s$ for all $p \int P_0$.
- $dp \perp 1, \mathrm{u}(b) - \mathrm{u}(a), \mathrm{X}(a) \Rightarrow dp \perp \mathrm{X}(b)$,
- Linear algebra
- $\Rightarrow \bar{\mathrm{X}}(b) \in \text{span}(\bar{\mathrm{X}}(a), \bar{\mathrm{u}}(b) - \bar{\mathrm{u}}(a))$,
- $=> X$ is aligned on $\{a, b\}$.

- Hence, for all $p \in P_0 = \{q : a \in A(q), r = \mathbb{E}_q X(a,.)\}$
  - (a, r) is optimal at P_0
  - the expected maximal scoring payoff is equal to the expected payoff from (a, r), hence affine over P_0,
  - take $q \in \text{int} P_0$ and let $s = \mathbb{E}_q X(b,.)$,
  - Observation $\Rightarrow (b, s)$ is optimal at all $p \int P_0$,
  - $\mathbb{E}_p X(b,.) = s$ for all $p \int P_0$.
- $dp \perp 1, \mathrm{u}(b) - \mathrm{u}(a), \mathrm{X}(a) \Rightarrow dp \perp \mathrm{X}(b)$,
- Linear algebra
- $\Rightarrow \bar{\mathrm{X}}(b) \in \mathrm{span}(\bar{\mathrm{X}}(a), \bar{\mathrm{u}}(b) - \bar{\mathrm{u}}(a))$,
- $=> X$ is aligned on $\{a, b\}$.

- Hence, for all $p \in P_0 = \{q : a \in A(q), r = \mathbb{E}_q X(a,.)\}$
  - (a, r) is optimal at P_0
  - the expected maximal scoring payoff is equal to the expected payoff from (a, r), hence affine over P_0,
  - take $q \in \text{int} P_0$ and let $s = \mathbb{E}_q X(b,.)$,
  - Observation $\Rightarrow (b, s)$ is optimal at all $p \int P_0$,
  - $\mathbb{E}_p X(b,.) = s$ for all $p \int P_0$.
- $dp \perp 1, \mathrm{u}(b) - \mathrm{u}(a), \mathrm{X}(a) \Rightarrow dp \perp \mathrm{X}(b)$,
- Linear algebra
- $\Rightarrow \bar{\mathrm{X}}(b) \in \text{span}(\bar{\mathrm{X}}(a), \bar{\mathrm{u}}(b) - \bar{\mathrm{u}}(a))$,
- $=> X$ is aligned on $\{a, b\}$.

Necessary conditions

Adjacency Lemma

- Adjacency Lemma => if X is incentivizable, then, for all adjacent actions a, b, there is x and y such that

$$\bar{\mathsf{X}}(a) = x\bar{\mathsf{X}}(a) + y\Delta_a^b.$$

Necessary conditions

Adjacency graph

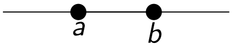

- Adjacency Lemma imposes restrictions on $\bar{X}(a)$ given $\bar{X}(b)$

$$\bar{X}(a) = x\bar{X}(b) + y\Delta_a^b$$

Necessary conditions

Adjacency graph

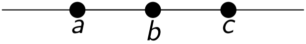

- These conditions carry over through adjacency paths ...

$$\begin{aligned}\bar{\mathsf{X}}(a) &= x\bar{\mathsf{X}}(b) + y\Delta_a^b \\ &= x'\bar{\mathsf{X}}(c) + y\Delta_a^b + y'\Delta_b^c\end{aligned}$$

Necessary conditions
Adjacency graph

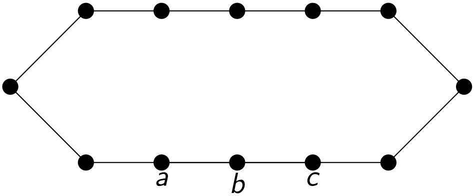

- ... and come back through cycles.

$$\bar{X}(a) = x_0 \bar{X}(a) + \sum_{b \in C \setminus \{a\}} y_b \Delta_a^b$$

Outline

Representation in special cases

	Tree	Complete graph	Product problem
Adjacency graph	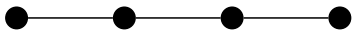	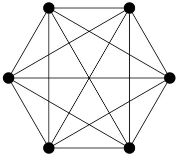	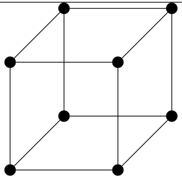
Examples	Cognitive uncertainty (Enke and Graeber, 2023), 1-dimensional supermodular problems,	Multiple choice question, prediction problems	Random problem selection (Azrieli et al., 2018), test with ≥ 3 questions
Representation	Piecewise aligned	Aligned	Product aligned

Special cases: Tree

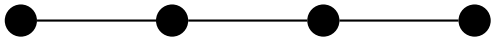

Example: (Enke and Graeber, 2023)

Special cases: Tree

- (Enke and Graeber, 2023): DM chooses certainty equivalent $a^*(p; q, y)$ of a lottery $O^{1-q}1^q$. The utility of the lottery is subject to cognitive uncertainty θ.
- CE is BDM incentivized:

$$a^*(p; q, y) = \arg\max_a \mathbb{E}_p \left[\frac{a}{y} q u_0(y, \theta) + \int_a^y p u_0(z, \theta) dz \right]$$
$$= \arg\max_a \mathbb{E}_p u(a, \theta)$$

- **Lemma**: The adjacency graph is a line.
- EG ask about the probability that the *ex post* correct CE is within ϵ of the chosen CE:

$$X(a, \theta) = \begin{cases} 1 & |a - a^*(\delta_\theta)| < \epsilon \\ 0 & \text{otherwise.} \end{cases}$$

- Adjacency Lemma => X is not incetivizable for generic u_0.

Special cases: Tree

Incentivizability

Theorem: Incentivizability on tree-like problems

Suppose that the adjacency graph is a tree.

Then, X is incentivizable if and only if it satisfies the Adjacency Lemma for each adjacent pair.

- Proof: scoring rules paste scoring rules over two disjoint set connected by a single adjacent pair.

Special cases: Complete graph

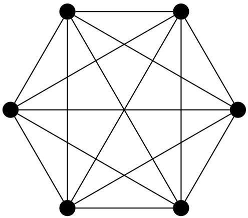

Example: multiple-choice question.

Theorem: Incentivizability on complete graphs

Suppose that $|A| \geq 4$, the adjacency graph is a complete, and for all actions a, b_0, b_1, b_2, vectors $\Delta_a^{b_0}, \Delta_a^{b_1}, \Delta_a^{b_2}$ are linearly independent.

Then, X is incentivizable if and only if it has aligned representation:

- There are $\gamma, \kappa : A \longrightarrow \mathbb{R}$, and $d \in \mathbb{R}^{\Theta}$ such that

$$\begin{aligned} X(a,\theta) =& \gamma(a)\left(u(a,\theta) + d(\theta)\right) + \kappa(a), or \\ =& \gamma(a)\left(d(\theta)\right) + \kappa(a) \end{aligned}$$

- Complete graphs have lots of cycles. Proof

Special cases: product problems

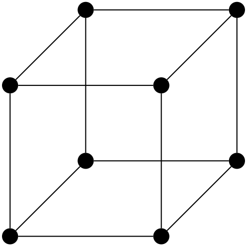

Example: Random problem selection

Special cases: product problems

Product problem: $(\Theta, A, u(.,.))$, where

- $(\Theta_i, A_i, u_i(.,.))$ is a collection of tasks,
- $\Theta = \times_i \Theta_i$, $A = \times_i A_i$, and

$$u(a,\theta) = \sum_i u_i(a_i, \theta_i),$$

- two actions $a, b \in A$ are adjacent if they differ in exactly one task: $a_{-i} = b_{-i}$ for some i
- => lattice structure of the adjacency graph,
- X depends on task i *trivially* if, for each a_{-i}, the vectors $\bar{\mathrm{X}}(a_i a_{-i})$ are collinear.

Special cases: product problems

Product-aligned representation

Question X is product aligned if there are parameters $\gamma(a), y_i, \kappa(a) \in \mathbb{R}$ and $d \in \mathbb{R}^\Theta$ such that for each a

$$X(a, \theta) = \gamma(a) \left(\sum_i y_i u_i(a_i, \theta) + d(\theta) \right) + \kappa(a)$$

- y_is do not have to be the same or all poisitive.

Special cases: product problems

Theorem: Incentivizability on in product games

Suppose that each task i is either

- binary ($|A_i| = 2$), or
- it has complete graph and vectors $\{\Delta_a^b, \Delta_a^c\}$ are linearly independent for all $a, b, c \in A_i$.

If X depends non-trivially on at least 3 problems, then it is incentivizable iff it is product-aligned.

- One needs at least 3 problems.

Special cases: product problems

Example: Quizz

- Möbius et al. (2022) A subject writes IQ test with multiple-choice questions.
- The test has two parts: visual and verbal, each part has $N > 2$ questions.
- The payoff is proportional to score: number of correct answers.

- "What is the difference between the two parts of the test?" corresponds to

$$X(a, \theta) = \sum_{i \leq N} 1\{a_i = \theta_i\} - \sum_{i > N} 1\{a_i = \theta_i\}$$

This question is incentivizable.

Special cases: product problems

Example: Quizz

- Möbius et al. (2022) A subject writes IQ test with multiple-choice questions.
- The test has two parts: visual and verbal, each part has $N > 2$ questions.
- The payoff is proportional to score: number of correct answers.

- "How likely your score is above 50%?" corresponds to

$$x(a,\theta) = \begin{cases} 1 & \sum_i 1\{a_i = \theta_i\} \geq N \\ 0 & \text{otherwise} \end{cases}.$$

THis question is NOT incentivizable.

Outline

Comments

$|\Theta| = 3$

- Adjacency Lemma does not bite when $|\Theta| = 3$.
- Limited practical interest (two action-independent questions to learn all beliefs).
- Possibly, the fact that the expected maximal payoff rom scoring rule is convex may lead to some information baout incentivizability.

Joint incentivizability

Questions $\mathrm{X}, \mathrm{Y} : A \longrightarrow \mathbb{R}^{\Theta}$ are *jointly incentivizable* if there exists $V : \mathbb{R}^2 \times A \times \Theta \longrightarrow [0, 1]$ st. for every $p \in \Delta(\Theta)$,

$$\arg\max_{a,r,s} \mathbb{E}_p V(r, s, a, \theta)$$
$$= \left\{ \left(\mathbb{E}_p X(a; \theta), \mathbb{E}_p Y(a; \theta), a \right) : a \in \arg\max_{b \in A} \mathbb{E}_p u(b; \cdot) \right\}.$$

Adjacency Lemma for 2 questions

Suppose that X and Y are jointly incentivizable. If actions a and b are adjacent, then there are ρ_X, ρ_Y and σ_x^y for $x, y = X, Y$, not all equal to 0, such that

$$\bar{\mathrm{X}}(b) = \rho_X (\bar{\mathrm{u}}(b) - \bar{\mathrm{u}}(a)) + \sigma_X^X \bar{\mathrm{X}}(a) + \sigma_X^Y \bar{\mathrm{Y}}(a)$$

$$\text{and} \qquad \bar{\mathrm{Y}}(b) = \rho_Y (\bar{\mathrm{u}}(b) - \bar{\mathrm{u}}(a)) + \sigma_Y^X \bar{\mathrm{X}}(a) + \sigma_Y^Y \bar{\mathrm{Y}}(a).$$

Comments
Multiple questions

Many questions

All $|\Theta| - 1$ questions are jointly incentivizable.

With $|\Theta| - 1$, we can ask about all beliefs.

Comments

Non-linear questions

- Our techniques only apply to linear questions.
- Lambert (2019) studies elicitation of “properties” of beliefs, where a property corresponds to a discrete or continuum partition of the simplex
- A simple necessary condition: elicitable property must have "convex inverse images".
- Example: variance is (action-independent) non-incentivizable.

Conclusions

- Sufficient conditions: Aligned questions (i.e., questions about affine transformations of payoffs) are incentivizable.
- Necessary conditions: Adjacency Lemma.
- **"Informal Theorem"** In three classes of decision problems, question X is incentivizable if and only if it satisfies the Adjacency Lemma.
- Special representations when the adjacency graph is complete or in product problems.
- Many other questions:
  - dynamic elicitation (signals?)
  - "robust" elicitation.

Theorem: Incentivizability on complete graphs

Suppose that $|A| \geq 4$, the adjacency graph is a complete, and for all actions a, b_0, b_1, b_2, vectors $\Delta_a^{b_0}, \Delta_a^{b_1}, \Delta_a^{b_2}$ are linearly independent.

Then, X is incentivizable if and only if it has aligned representation:

Lemma 1

Suppose $C = (a_0, ..., a_{n-1})$ is a cycle such that vectors $\Delta_{a_0}^{a_1}, ..., \Delta_{a_0}^{a_{n-1}}$ are linearly independent.
Then, if X is incentivizable, then it is either aligned on C, or $\bar{X}(a) \in \text{span}\{\Delta_{a_0}^{a_1}, ..., \Delta_{a_0}^{a_{n-1}}\}$.

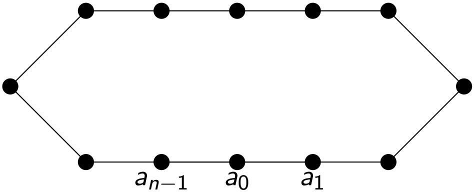

- Iteration of Adjacency Lemma => there exist x and y_i st.

$$\bar{X}(a_0) = x\bar{X}(a_0) + \sum_i y_i \Delta_{a_i}^{a_{i+1}}$$

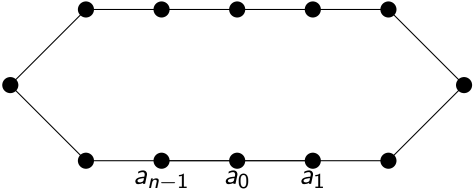

- Because $\Delta_{a_i}^{a_{i+1}} = \Delta_{a_0}^{a_{i+1}} - \Delta_{a_0}^{a_i}$, we have

$$\bar{X}(a_0) = x\bar{X}(a_0) + \sum_i (y_{i-1} - y_i)\Delta_{a_0}^{a_i}$$

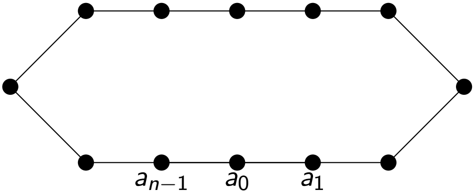

- If $X \notin \text{span}\{\Delta_{a_0}^{a_1}, ..., \Delta_{a_0}^{a_{n-1}}\}$, linear independence implies that $x = 1$ and $y_{i-1} = y_i$:

$$\bar{X}(a_0) = x\bar{X}(a_0) + \sum_i (y_{i-1} - y_i)\Delta_{a_0}^{a_i}$$

, which implies aligned representation on C.

Complete adjacency graph

Proof

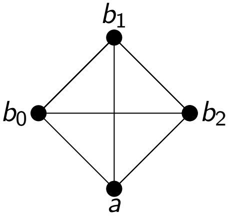

- Fix a and consider 3-cycles.

Complete adjacency graph
Proof

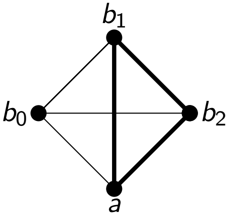

- Fix a and consider 3-cycles.

Complete adjacency graph

Proof

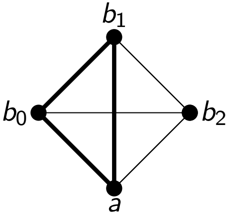

- The intersection of the spans of vectors for each cycle due to linear independence.

Complete adjacency graph

Proof

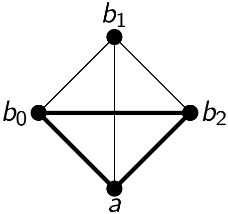

- $\bar{X}(a)$ cannot belong to all of them.

Complete adjacency graph

Proof

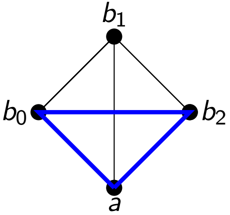

- So, there must be a cycle with aligned representation.

Complete adjacency graph

Proof

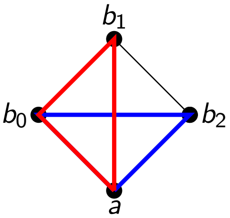

- We can apply the same argument to the remaining action.
- But, the two "alignments" do not have to be the same.

Lemma 2 (merging representations)

Suppose X is aligned on B and C and $a, b \in B$, $a \neq b$ are such that $\bar{X}(a)$ and $\bar{X}(b)$ are not collinear.
Then, X is aligned on $B \cup C$.

Go back