Reaping the Informational Surplus in Bayesian Persuasion

Rann Smorodinsky

Joint work with

Niklas Hahn, Martin Hoefer, and Ronen Gradwohl

Motivation

Policymakers often turn to experts for advice on the potential impacts of a policy.

Experts may have different incentives from policymaker. For example:

- Policymaker cares about impact on society as a whole;
- Experts care about impact on society, but also about impact on particular sector.

Aumann and Maschler (1966) and Kamenica and Gentzkow (2011) show how experts can enjoy their informational advantage.

But ... what if there is more than one expert for the policymaker to choose from?

Outline

- One long example
- Two surprising observations
- Model and results
- Related literature

Example

A policymaker must decide whether to adopt a policy that restricts e-cigarettes (P) or maintain the status-quo (Q).

The policy either improves public health (H) or not ($\overline{H}$).
The policymaker prefers P iff it improves public health.

In addition, the policy will either cause a decrease in tobacco-use by youth (Y) or an increase ($\overline{Y}$).

There is an expert E who

- Similar to policymaker, cares about public health;
- also cares about decrease in tobacco-use by youth.

State space = $\{H, \overline{H}\} \times \{Y, \overline{Y}\}$.

- Policymaker prefers P in states (H, Y) and $(H, \overline{Y})$;
- Expert prefers P in states (H, Y), $(H, \overline{Y})$ and $(\overline{H}, Y)$.
- Utilities are 1 when preferred action is taken, and 0 otherwise.

Let $\Pr[H] = \alpha \geq 0.5$ and $\Pr[Y] = \beta \in (0, 1)$, independently.

Claim: E's optimal scheme is one that always recommends the action that is optimal to him.

- Namely, recommend Q in state $(\overline{H}, \overline{Y})$, otherwise recommend P.

Under this scheme, policymaker wants to follow recommendations.

Note that outcomes are optimal for expert, not policymaker.

And what if there are two experts?

Now suppose there are two experts, E_1 and E_2.
Each cares about public health, but

- E_1 also cares about tobacco-use by youth, and wants it to decrease (Y);
- E_2 also cares about tobacco-use by women, and wants it to decrease (W).

State space is $\{H, \overline{H}\} \times \{Y, \overline{Y}\} \times \{W, \overline{W}\}$.

E_1 wants policy iff *either* H or Y are realized.

- That is, in every state except $(\overline{H}, \overline{Y}, W)$ or $(\overline{H}, \overline{Y}, \overline{W})$.

E_2 wants policy iff *either* H or W are realized.

- That is, in every state except $(\overline{H}, Y, \overline{W})$ or $(\overline{H}, \overline{Y}, \overline{W})$.

Utilities are 1 when preferred action is taken, and 0 otherwise.

$\Pr[H] = \alpha \geq 0.5$ and $\Pr[Y] = \Pr[W] = \beta \in (0, 1)$, all independently.

Experts' signal can only depend on their payoff relevant states:

- E_1 about H and Y, and E_2 about H and W.

The two experts game

- Both experts propose a signaling scheme;
- Policymaker chooses an expert;
- State is realized;
- Policymaker chooses an action based on chosen expert's signal.

Theorem: In every (mixed) equilibrium, policymaker always takes his correct action.

Policymaker gets utility 1 and fully reaps the informational surplus. Experts' utility is $\alpha + (1-\alpha)(1-\beta)$.

Intuition

Suppose chosen expert (say, E_1) does not always recommend policymaker-optimal action. Obviously, E_1 must recommend P whenever the state is H. So, it must be the case that E_1 recommends P with probability $\delta > 0$ whenever the state $\overline{H}$ holds (e.g., when Y holds).

But then E_2 has a profitable deviation:

- Simulate E_1's signal - In state H signal P and in state $\overline{H}$ signal P with probability δ.
  This will make the policy make indifferent between the two experts.
- Correlate the latter with the event W.
  This will make E_2 better off, if chosen.
- Slightly reduce the probability of P in state $\overline{H}$ at the cost of "shaving" E_2's utility.
  Now the policymaker strictly prefers E_2 to E_1.

Implications

- Policymaker's point of view: the restriction to interacting with just a single expert is actually a benefit, and in fact should be self-imposed.
- Experts' point of view: commitment power could be a double-edged sword.

The benefit of restricting to a single expert

What if Policymaker were to observe *both* experts' signal realizations?

Assume for simplicity that $\alpha = 0.5$.

Consider the strategy tuple where both experts adopt their expert-optimal schemes from the single-expert setting and, in reply, Policymaker chooses P iff *both* experts recommend P.

Observation: This is an equilibrium (not a unique one).

Policymaker's utility is $\frac{1}{2} + \frac{1}{2}(1 - \beta^2) < 1$.

The double-edged sword of commitment power

In the single-expert setting, commitment power is always beneficial. Is this statement correct when there are many experts?

With commitment power: expert utilities are $1/2 + (1-\beta)/2$.

And what if they have no commitment power?

- Interaction between chosen expert and policymaker is cheap talk: expert must recommend own preferred action, and policymaker obliges.
- Policymaker is indifferent between experts, so chooses at random.
- Chosen expert's utility is 1.
- Non-chosen expert's utility is $1 - \beta(1-\beta)$.
  (Non-chosen expert gets utility 1 in all cases except when the policy is bad for the policymaker and the two experts disagree.)

Experts' ex-ante expected utility is $\frac{1+1-\beta(1-\beta)}{2} > 1/2 + (1-\beta)/2$.

The model

n senders denoted $\{1, \ldots, n\}$ and a receiver denoted R.

Finite state space $\Omega = \Omega_1 \times \ldots \times \Omega_n \times \Omega_R$, with common prior μ.

Each sender i will only learn realization of Ω_i and Ω_R.

Receiver has action set $\mathcal{A}$.

Payoff functions $u_R : \mathcal{A} \times \Omega_R \mapsto \mathbb{R}$ and $u_i : \mathcal{A} \times \Omega_i \mapsto \mathbb{R}$ for each i.

Generalization* of standard model where all players have common state space Ω.

Simplifying assumption: For every $j \in \{1, \ldots, n, R\}$ and $\omega_j \in \Omega_j$, the set $\arg\max_{a \in \mathcal{A}} u_j(a, \omega_j)$ is a singleton.

Main assumption

For every pair of senders $i \neq j$ and every $(\omega_j, \omega_R) \in \Omega_j \times \Omega_R$ in $\text{supp}(\mu)$, there exists an $\omega_i \in \Omega_i$ with $P(\omega_i \mid (\omega_j, \omega_R)) > 0$ for which $\arg\max_{a \in \mathcal{A}} u_i(a, \omega_i) = \arg\max_{a \in \mathcal{A}} u_R(a, \omega_R)$.

- Holds in e-cigarette example.
- Holds whenever μ has full support, and there are no undesirable actions—for each sender i and each action a, there is some state ω_i for which $a = \arg\max_{a' \in \mathcal{A}} u_i(a', \omega_i)$.
- Holds for any "perturbed" prior.

The game

A signaling scheme for sender i is a function $\pi_i : \Omega_i \times \Omega_R \mapsto \Delta(\mathcal{A})$.

Let Π_i denote all signaling schemes of player i.

1. Each sender i chooses a distribution $\tilde{\pi}_i$ over Π_i.
2. Receiver observes the **realized** vector $(\pi_1, \ldots, \pi_n)$.
3. Receiver chooses a player, say j (breaking ties with a coin flip).
4. State is realized and Receiver observes (only) $\pi_j(\omega_{jR})$.
5. Receiver chooses an (optimal) action from $\mathcal{A}$.

Equilibrium

- $v_R(\pi)$ is the receiver's expected utility.
- $v_i(\pi)$ is the sender i's expected utility.

A *NE* is a profile $\tilde{\pi} = (\tilde{\pi}_1, \ldots, \tilde{\pi}_n)$ with the following property: There do not exist a sender i and distribution $\tilde{\pi}'_i$ over signals with

$$E\left[v_i(\tilde{\pi}'_{D(\tilde{\pi}')})\right] > E\left[v_i(\tilde{\pi}_{D(\tilde{\pi})})\right],$$

- where $\tilde{\pi}'$ is the profile $\tilde{\pi}$ but with $\tilde{\pi}'_i$ replacing $\tilde{\pi}_i$,
- and $D(\pi)$ returns a uniformly random index from the set $\{j : v_R(\pi_j) \geq v_R(\pi_i)\ \forall i \in \{1, \ldots, n\}\}$.

Fully-informative equilibrium

Definition: π_i is *fully informative* if, whenever an action is recommended by sender i, that action is optimal for the receiver.

Definition: A profile $\tilde{\pi}$ is *fully informative* if, with probability 1, the profile $(\pi_1, \ldots, \pi_n)$ drawn from $\tilde{\pi}$ satisfies the following: For every sender j that is chosen by the receiver with positive probability it holds that π_j is fully informative.

Our main result: Every NE of the game is fully informative.

Main idea: Suppose NE is the pure profile $(\pi_1, \ldots, \pi_n)$.

Suppose player j is always chosen, but that π_j is not fully informative.

We construct a profitable deviation for any other player i.

First, i *simulates* the signaling scheme of player j, by playing a scheme π_i' that yields the same marginal distribution over $\Omega_i \times \Omega_R$ as π_j.

- Both i and R are are indifferent between π_j and π_i'.

Second, i modifies π_i' to yield a scheme π_i'' that both he and R strictly prefer to π_i' and hence to π_j.

- By making π_i' a touch more informative (to improve the receiver's utility), but only when his own preferences are aligned with the receiver's (to also improve his own utility).

What if j chosen with positive probability < 1?

Let i be another sender chosen with positive probability.

If i weakly prefers π_j to π_i, then proceed as before, yielding profitable deviation to i.

If i strictly prefers π_i to π_j, then i modifies π_i to be fully-revealing with probability $\varepsilon > 0$.

Now i is chosen with probability 1.

Profitable when ε sufficiently small.

Related literature

Bayesian persuasion was initiated by Aumann and Maschler (1966), recently popularized by Kamenica and Gentzkow (2011).

Most closely related to work featuring competition between senders.

- Gentzkow and Kamenica (2017a,b): When does competition increase informativeness? Blackwell-connectedness. Li and Norman (2018).
- Au and Kawai (2020): Senders whose only aim is to be "chosen" by receiver.
- Board and Lu (2018), Au (2018), and Whitmeyer (2020): Sequential search with endogenous prize distributions. Lower search costs decrease informativeness.
- Boleslavsky and Cotton (2018): Restricting action space may increase informativeness.

Summary

A receiver that can choose which sender to work with can reap all the informational surplus.

1. Receivers are better off restricting attention to a single sender.

2. Senders may be better off forgoing their commitment power.

Thank You!