Lecture 4: Cheap Talk

Alexander Wolitzky

MIT

14.126, Spring 2024

Cheap Talk

What is the effect of costless communication in strategic interactions with asymmetric information?

- Formally, cheap talk games are signaling games where the signal ("message") is payoff-irrelevant. Typically consider what's achievable with arbitrary messages, unlike signaling games where the signal set is fixed.
- In any cheap talk game, there's a PBE where messages are uninformative: a "babbling equilibrium."
- In some games, all PBE are uninformative: e.g., game is 0-sum, or sender always wants receiver to take $a = 1$ rather than 0.
- Informative PBE can exist if there is some "alignment" between sender's and receiver's preferences. A theme of cheap talk models is that some alignment exists more often than one might think.

Plan

1. Canonical model with 1 sender, 1 receiver, 1-dimensional state and action. (Crawford Sobel 82)
2. Multidimensional state + action (Chakraborty Harbaugh 10, Lipnowski Ravid 20)
3. Multiple senders (Battaglini 02, Ambrus Takahashi 08)
4. Mutiple receivers (Farrell Gibbons 89)
5. Multiple rounds of communication (Aumann Hart 03, Krishna Morgan 04)
6. Mediated communication (Blume Board Kawamura 08, Goltsman et al 09)

Cheap Talk: General Framework

- Sender observes state $\theta \in \Theta$, drawn from prior $\mu \in \Delta(\Theta)$.
- Sender chooses a message $m \in M$. Typically ask what can be achieved with "sufficiently rich" M. (Next slide.)
- Receiver observes m but not θ, takes an action $a \in A$.
- Preferences $u_S(a, \theta)$ for sender, $u_R(a, \theta)$ for receiver.

How Rich is Rich Enough?

By an argument akin to the revelation principle, it is always without loss to take $M = \Delta(A)$.

- Fix any message set M and any PBE.
- In this PBE, each message m induces some distribution over receiver actions $\alpha_m \in \Delta(A)$.
- Now replace M with $\Delta(A)$. Have sender send α_m whenever she sent m in the original PBE. Have receiver take α_m following each on-path message α_m, take an arbitrary fixed α_{m_0} following any off-path message α.
- It's optimal for receiver to take α_m when sender says α_m, as his belief is the same as it was following m in the original PBE. (Can interpret off-path messages α as if they were α_{m_0}.)
- It's optimal for sender to send α_m whenever she sent m in the original PBE, as the set of mixed receiver actions she can induce is weakly smaller.

In many cheap talk settings, can take M much smaller than this, such as $M = A$ or $M = \Theta$.

Comparison with Other Communication Models

Receiver cannot commit to a decision rule (map from messages to actions). This makes cheap talk different from screening or delegation. (Holmström 84, Alonso Matouschek 08)

- More outcomes are implementable under delegation, because delegation eliminates receiver's IC constraint.

Sender cannot commit to a disclosure policy (map from states to distribution over messages). This makes cheap talk different from Bayesian persuasion. (Kamenica Gentzkow 11)

- More outcomes are implementable under Bayesian persuasion, because BP eliminates sender's IC constraint.

Set of feasible messages does not depend on the state. This makes cheap talk different from models with verifiable ("hard") information. (Grossman 81, Milgrom 81)

Other Communication Models (cntd.)

One can think of many other variants that affect commitment in some way. E.g., under *mediated communication*, Sender commits in advance to the design of a noisy communication channel, but then chooses what to input to the channel. (Or R designs channel.)

Yet more variants arise if Receiver also has private information (e.g., can Sender elicit it before communicating her own information?), if Sender's private information is endogenous (e.g., when does Receiver benefit form Sender being well-informed? when does Sender benefit from Sender being well-informed, if Receiver knows Sender's info structure?), or if communication is combined with other incentive instruments (e.g., transfers, fines, delegation).

Due to revival of interest in communication following Bayesian persuasion, and fact that the Bayesian persuasion commitment assumption is often inappropriate, there is much current interest in different models of communication under various assumptions.

PBE

PBE definition is as in signaling games:

Definition

A ***p*erfect Bayesian equilibrium (PBE)** consists of a strategy $\sigma \in \Delta(M)^{\Theta}$ for the sender, a strategy $\alpha \in A^M$ for the receiver, and a belief system $\mu \in \Delta(\Theta)^M$ for the receiver, such that

1. Sender IC: For all θ,
   $\sigma(m|\theta) > 0 \implies m \in \text{argmax}_{m \in M} u_S(\theta, \alpha(m))$.
2. Receiver IC: For all m,
   $\alpha(a|m) > 0 \implies a \in \text{argmax}_{a \in A} u_R(\mu(m), a)$.
3. Bayes' consistency: For all on-path m, $\mu(m)$ is given by Bayes' rule.

Babbling, Fully Revealing, Partially Revealing

In a **babbling equilibrium**, $\sigma(\cdot|\theta)$ is independent of θ, and the receiver takes a best response given the prior.

- A babbling equilibrium always exists.
- A non-babbling equilibrium is called **informative**.

In a **fully revealing equilibrium**, the supports of $\sigma(\cdot|\theta)$ and $\sigma(\cdot|\theta')$ are disjoint for all $\theta \neq \theta'$, and the receiver takes a best response $a_R(\theta)$ given the true state.

- If receiver's BR $\operatorname{argmax}_{a \in A} u_R(a, \theta)$ is unique for each θ, a fully revealing equilibrium exists iff $u_S(a_R(\theta), \theta) \geq u_S(a_R(\theta'), \theta)$ for all θ, θ'.

An equilibrium that is neither babbling nor fully revealing is called **partially revealing**.

Babbling / fully revealing / partially revealing equilibria are analogous to pooling / separating / semi-separating equilibria in signaling, but in cheap talk emphasize partially revealing equilibria.

Some Special Cases

Common interests: $u_S(a, \theta) = u_R(a, \theta)$ for all a, θ. (More generally, S and R have same ordinal preferences over a for each θ.) A fully revealing equilibrium exists.

Opposing interests (0-sum game): $u_S(a, \theta) = -u_R(a, \theta)$ for all a, θ. All equilibria are payoff-equivalent to babbling.

- Does *not* extend to the case where S and R have opposite ordinal preferences but not opposite cardinal preferences.
- Schelling 1960 already noted the difference in cheap talk's impact with common vs. opposed interests. Thinking about intermediate cases was the impetus for Crawford Sobel 1982.

State-independent sender preferences: $u_S(a, \theta) = u_S(a, \theta')$ for all a, θ, θ'. An informative equilibrium may exist, but in every equilibrium S is always indifferent over all messages that are ever sent with positive probability.

Special Cases (cntd.)

Quadratic preferences: $\Theta, A \subset \mathbb{R}$, and there is some "bias" $b > 0$ such that

$$u_R(a,\theta) = -(\theta - a)^2, \qquad u_S(a,\theta) = -(\theta + b - a)^2.$$

Both parties' utilities determined by the expected residual variance:

$$\begin{aligned}\mathbb{E}[u_R(a,\theta)] &= -\mathbb{E}_m[\mathrm{Var}(\theta|m)],\\ \mathbb{E}[u_S(a,\theta)] &= -\mathbb{E}_m[\mathrm{Var}(\theta|m)] - b^2.\end{aligned}$$

- With quadratic preferences, regardless of b, S and R have the same preferences *over equilibria*: both prefer more informative equilibria, in the sense of lower expected residual variance.

Finite-State Example

Consider quadratic preferences with $\Theta = \{0, 1/4, 1/2, 3/4\}$, $A = [0, 1]$, full-support prior, $b = 3/16$.

There is no fully revealing equilibrium.

- If there were, R would take $a = \theta$ for each θ.
- But then type $\theta = 0$ S could send message of type $\theta = 1/4$. This gives a loss of $(1/16)^2$ instead of $(3/16)^2$, so it's a profitable deviation.

Finite-State Example (cntd.)

Now consider a partially revealing **"interval equilibrium,"** where types 0 and 1/4 send one message (call it m_1), and types 1/2 and 3/4 send another messsage $m_2 \neq m_1$.

- Sender IC requires that type 1/4 prefers $a_1 = \mathbb{E}[\theta|\theta \in \{0, 1/4\}]$ to $a_2 = \mathbb{E}[\theta|\theta \in \{1/2, 3/4\}]$.
- (This implies that type 0 also prefers a_1 to a_2. Types 1/2 and 3/4 always prefer a_2 to a_1, since $b > 0$.)
- Type 1/4's prefers whichever of a_1 and a_2 is closer to 7/16.
- This is indeed a_1 if the prior is s.t. 1/4 is likely compared to 0 (so a_1 is close to 1/4) and 3/4 is likely compared to 1/2 (so a_2 is close to 3/4). If instead 1/4 and 3/4 are unlikely, there is no informative interval equilibrium.

Alternatively, suppose the states are equally likely and vary $b > 0$.

- A fully revealing equilibrium exists if b is sufficiently small.
- This relies on discrete states, so S can't shade "just a bit."
- If fix b and take the grid finer, eventually break the fully revealing equilibrium.

Continuous State and Actions (CS 82)

CS 82 consider $\Theta = A = [0, 1]$ and a class of preferences that generalizes quadratic.

- u_S, u_R are twice continuously differentiable, strictly concave in a, and satisfy strict single-crossing in (a, θ): for $a < a'$, $\theta < \theta'$, $i \in \{S, R\}$, we have $u_i(a, \theta) = u_i(a', \theta) \implies u_i(a, \theta') < u_i(a', \theta')$.
- For each θ, let $a_S(\theta), a_R(\theta)$ denote the preferred actions of S, R given θ. (Given above assumptions, these are uniquely defined, continuous, and increasing.) Also, let $a_R(\tilde{\mu})$ be R's optimal action at belief $\tilde{\mu} \in \Delta(\Theta)$.
- Assume that there exists $\varepsilon > 0$ such that $a_S(\theta) > a_R(\theta) + \varepsilon$ for all θ.

Interval Partition Equilibria

For each $\theta' \leq \theta''$, let

$$a_R\left(\theta', \theta''\right) = \underset{a \in A}{\operatorname{argmax}}\, \mathbb{E}\left[u_R\left(a, \theta\right) | \theta \in \left[\theta', \theta''\right]\right].$$

- a_R is continuous and strictly increasing in each argument.

A triple $\left(\sigma, a_R\left(\cdot, \cdot\right), \mu\right)$ is an **interval partition equilibrium** if there exists a partition of $[0, 1]$ into intervals $\left\{\left[\theta_{n-1}, \theta_n\right]\right\}_{n=1}^{N}$ with $\theta_0 = 0$ and $\theta_N = 1$ such that, for all $n \in \{1, \ldots, N\}$, if $\theta \in \left[\theta_n, \theta_{n+1}\right]$ then

1. Each message of S conditional on θ reveals that the state lies in $\left[\theta_n, \theta_{n+1}\right]$ (and no further information).
2. Each action of R conditional on θ equals $a_R\left(\theta_n, \theta_{n+1}\right)$.

All PBE are Interval Partitional

Lemma

Every PBE is an interval partition equilibrium, and the number of intervals in any such equilibrium is bounded.

Proof. Fix an arbitrary PBE.

- For each $a \in [0,1]$, let $M(a) = \{m \in M : \alpha(m) = a\}$. (Possibly empty. Note that $\alpha(m) = \mathbb{E}[\theta|m]$, and hence is pure.)
- Say that θ *induces* a if θ assigns prob >0 to a message in $M(a)$.
- Let $\hat{A}$ be the set of actions induced by any type. Note that if θ induces a then $u_S(a,\theta) \geq u_S(a',\theta)$ for all $a' \in \hat{A}$.
- Since u_S is strictly concave in a, each type θ can induce at most two actions.

Proof (cntd.)

We show that $a' - a > \varepsilon$ for all $a, a' \in \hat{A}$ s.t. $a < a'$.
(Recall: $a_S(\theta) > a_R(\theta) + \varepsilon \ \forall \theta$.)

- Note that there exists θ^* s.t. $u_S(a, \theta^*) = u_S(a', \theta^*)$.
  This follows because there exist $\theta \leq \theta'$ s.t.
  $u_S(a, \theta) \geq u_S(a', \theta)$ and $u_S(a, \theta') \leq u_S(a', \theta')$, so by continuity there exists $\theta^* \in [\theta, \theta']$ where they're equal.
- By strict single crossing of u_S, $a \leq a_S(\theta^*) \leq a'$, and each type $\theta < \theta^*$ strictly prefers a to a', and vice versa for $\theta > \theta^*$.
- So, R's beliefs have support $\Theta(a) \subset [0, \theta^*]$ when he takes a, and support $\Theta(a') \subset [\theta^*, 1]$ when he takes a'.
- By single crossing of u_R, this implies that $a \leq a_R(\theta^*) \leq a'$.
- We have $a_S(\theta^*) \in [a, a']$, $a_R(\theta^*) \in [a, a']$, and $a_S(\theta^*) - a_R(\theta^*) > \varepsilon$. So, $a' - a > \varepsilon$.

In particular, $|\hat{A}| < 1/\varepsilon$.

Proof (cntd.)

We showed that for all $a < a'$, there exists θ^* s.t. $\Theta(a) \subset [0, \theta^*]$ and $\Theta(a') \subset [\theta^*, 1]$.

- The convex hulls of the sets of states that induce distinct actions do not overlap (except at a single boundary point).
- Since each state induces some action, the supports of states that induce distinct actions must be disjoint intervals, which partion $[0, 1]$.

And, since $\left|\hat{A}\right| < 1/\varepsilon$, there are at most $1/\varepsilon$ such intervals.

Equilibrium Characterization

Theorem

A partition $\{[\theta_{n-1}, \theta_n]\}_{n=1}^{N}$ *constitutes a PBE if and only if* $\theta_0 = 0$, $\theta_N = 1$, *and for every* $n \in \{1, \ldots, N-1\}$,

$$u_S\left(a_R\left(\theta_{n-1}, \theta_n\right), \theta_n\right) = u_S\left(a_R\left(\theta_n, \theta_{n+1}\right), \theta_n\right).$$

In addition, there exists $\bar{N} \geq 1$ *such that there exists a PBE with* N *intervals if and only if* $1 \leq N \leq \bar{N}$.

Remarks

- Continuum of IC constraints for S and R reduces to indifference for cutoff sender types and $a = a_R(\theta_n, \theta_{n+1})$. Indifference for cutoff sender type suffices for sender IC by strict single-crossing.
- Indifference condition together with $\theta_0 = 0$, $\theta_N = 1$ determines $\{\theta_1, \ldots, \theta_{N-1}\}$ as a solution to a 2nd-order difference equation.
- There can be multiple solutions for the same N. The last part of the theorem follows because whenever a solution exists for N intervals, another solution exists for $N - 1$ intervals. (Original CS proof of this actually contains an error. Corrected by Kandori Kono 19.)
- In the quadratic preferences case, S and R have aligned preferences over PBE, with lower-residual variance equilibria better for both. In this case higher-N equilibria Pareto dominate lower-N equilibria.

Quadratic Preferences Example

Consider quadratic preferences with $\theta \sim \mathit{Uniform}\,[0,1]$.

With quadratic preferences, $a_R = \mathbb{E}\left[\theta | m\right]$. So, with quadratic preferences + uniform state, $a_R\left(\theta_n, \theta_{n+1}\right) = \left(\theta_n + \theta_{n+1}\right)/2$.

So, cutoff sender type indifference becomes

$$\left(\frac{\theta_{n-1}+\theta_n}{2} - \theta_n - b\right)^2 = \left(\frac{\theta_n+\theta_{n+1}}{2} - \theta_n - b\right)^2,$$

which simplifies to

$$a_{n+1} = 2a_n - a_{n-1} + 4b.$$

This difference equation implies that the maximum number of intervals $\bar{N}$ is decreasing in b.

- This yields a key take-home message of CS 82: more aligned preferences $\implies$ more scope for informative communication.
- An important goal of the cheap talk literature is understanding how this messages generalizes to richer environments.

Equilibrium Selection

CS 82 informally argued that it's natural to focus on the most informative (highest N) equilibrium.

Chen Kartik Sobel 08 provide some formal support for this.

Say that a PBE satisfies **no incentive to separate (NITS)** if type 0 sender would not benefit from revealing her type (if she could): that is, $u_S\left(a_R\left(0,\theta_1\right)0\right) \geq u_S\left(a_R\left(0\right)0\right)$.

Intuition: since Sender is biased up, seems natural that she can credibly "confess" to having the lowest possible type.

CKS show that every equilibrium with $\bar{N}$ intervals satisfies NITS, and that under an additional condition (which holds in the uniform-quadratic case) there is a unique equilibrium with N intervals for each $N \leq \bar{N}$, and **only** the equilibrium with $\bar{N}$ intervals satisfies NITS.

Equilibrium Selection (cntd.)

Moreover, when $M = [0, 1]$ (so messages can have "literal meanings"), CKS show that NITS holds in every equilibrium that satisfies two conditions:

1. **Monotonicity:** $m(\theta)$ and $a(m)$ are both non-decreasing.
2. **"Grain of honesty/gain of credulity"**: Either one of the following holds:
   2.1 There exist probabilities $\sigma, \rho > 0$ such that with probability σ Sender non-strategically takes $m(\theta) = \theta \; \forall \theta$, and with independent probability ρ Receiver non-strategically takes $a(m) = m \; \forall m$, and we take the limit as $\sigma, \rho \to 0$.
   2.2 Sender's preferences include a "lying cost": sender's utility is $u_S(a, \theta) - kC(m, \theta)$ for $k > 0$, C continuous, $\partial C / \partial m < 0$ if $m < \theta$, $\partial C / \partial m > 0$ if $m > \theta$, and we take the limit as $k \to 0$.

To what extent one can select NITS / most informative equilibrium using only monotonicity **or** "grain of honesty"-type assumptions (rather than both together) is not clear.

Multidimensional State + Action

From the perspective of general communication games, CS's 1-dimensional state + action model is quite special.

Another tractable model is to allow multidimensional Θ and A, but assume sender's preferences are **state-independent:** $u_S(a,\theta) = u_S(a,\theta')$ $\forall a, \theta, \theta'$.

- E.g. sender just wants a high wage, or just wants receiver to accept her proposal.
- This model is studied by Chakraborty Harbaugh 07, 10; Lipnowski Ravid 20.
- Key idea: sender must get same payoff from all equilibrium messages, but this still leaves room for informative communication.

Chakraborty Harbaugh: Example

- Sender is a seller with 2 goods to sell to a buyer (Receiver), who can buy at most 1.
- Sender gets payoff of 1 if Receiver buys, 0 otherwise.
- Good i has value $v_i \sim U[0,1]$ for buyer, independent across goods. Sender knows (v_1, v_2).
- Buyer's utility of not buying is $\varepsilon \sim U[0,1]$. Buyer knows ε.
- Is cheap talk valuable?
- Yes: suppose sender reports **which** v_i is higher.
- Then buyer buys good with higher v_i with prob $\mathbb{E}[\max\{v_1, v_2\}] = 2/3$, vs. $1/2$ without cheap talk.
- Both parties are better off than without cheap talk.

Chakraborty Harbaugh: Main Result

Main result: informative equilibria exist very generally when $\dim \Theta \geq 2$ and $a_R = \mathbb{E}[\theta|m]$ (quadratic Receiver preferences).

Theorem

Assume that $\Theta \subset \mathbb{R}^d$ with $d \geq 2$ and Θ is compact, convex, and has non-empty interior; the prior μ has a full support density; $a_R = \mathbb{E}[\theta|m]$; and u_S is state-independent and continuous. Then an informative equilibrium exists.

Idea: communicate about a direction orthogonal to Sender's bias.

Proof

- Fix any point $\theta_0 \in \text{int}(\Theta)$. For any unit vector $\lambda \in \mathbb{R}^d$, let H_λ be the hyperplane through θ_0 with normal vector λ.
- H_λ partitions Θ into two positive-probability sets, $\Theta_\lambda^- = \{\theta : \lambda \cdot \theta < \lambda \cdot \theta_0\}$ and $\Theta_\lambda^+ = \{\theta : \lambda \cdot \theta > \lambda \cdot \theta_0\}$ (excluding the 0-prob boundary).
- Let $a_\lambda^- = \mathbb{E}\left[\theta | \theta \in \Theta_\lambda^-\right]$ and $a_\lambda^+ = \max_{a \in A} \mathbb{E}\left[\theta | \theta \in \Theta_\lambda^+\right]$. Note that these are continuous in λ.
- Let $\Delta_S(\lambda) = u_S\left(a_\lambda^+\right) - u_S\left(a_\lambda^-\right)$. If we can find λ such that $\Delta_S(\lambda) = 0$, we will have found an informative equilibrium.
  - Note: this conclusion uses state-independent sender preferences.

Proof (cntd.)

- Note that $\Delta_S(\lambda)$ is a continuous function on the unit sphere.
- By the Borsuk-Ulam theorem, there exists λ^* such that $\Delta_S(\lambda^*) = \Delta_S(-\lambda^*)$.
- (Borsuk-Ulam: for any continuous function $f : S^n \to \mathbb{R}^n$, there exists $x \in S^n$ such that $f(x) = f(-x)$.)
- Note also that, for all λ, $a_\lambda^+ = a_{-\lambda}^-$, and hence $\Delta_S(-\lambda) = -\Delta_S(\lambda)$.
- Hence, $\Delta_S(\lambda^*) = -\Delta_S(\lambda^*) = 0$.

Ranking Equilibria

Clearly, Receiver prefers any informative equilibrium to the babbling equilibrium.

So does Sender, when her preferences are quasi-convex.

- If actions $a_1, \ldots, a_N$ are induced in an informative equilibrium, by Receiver IC the action a_0 induced in the babbling equilibrium satisfies

$$a_0 = \mathbb{E}[\theta] = \mathbb{E}[\mathbb{E}[\theta|m]] = \mathbb{E}[a(m)],$$

  where $a(m) \in \{a_1, \ldots, a_N\}$ is the action taken at message m in the informative equilibrium. In particular, a_0 lies in the convex hull of $a_1, \ldots, a_N$.
- By quasi-convexity, $u_S(a_0) \leq \max_n u_S(a_n)$.
- By Sender IC, $u_S(a_1) = \ldots = u_S(a_N)$.
- So Sender prefers the informative eqm.

Ranking Equilibria: Remarks

This argument uses Sender indifference. If we only knew that $a_0 = \mathbb{E}[a(m)]$, we could conclude that Sender prefers the informative equilibrium only if her preferences are convex, not just quasi-convex.

Punchline: if Sender's preferences are quasi-convex, then both parties prefer more informative equilibria, as in CS.

Conversely, Sender prefers the babbling equilibrium if preferences are quasi-concave.

- Not clear if informative equilibria are reasonable in this case, as they violate Sender's "right to remain silent."
  However, standard refinements don't eliminate them.

Lipnowski Ravid

Lipnowski Ravid provide a geometric characterization of Sender's best equilibrium payoff in Chakraborty Harbaugh's model with general receiver preferences.

Lipnowski Ravid: Example

$\theta \in \{1, 2\}$, $\mu(1) = \mu(2) = \frac{1}{2}$, $a \in \{0, 1, 2\}$, $u_S(a) = a$,
$u_R(0, \theta) = 0$, $u_R(a \neq 0, \theta \neq a) = -1$, $u_R(a \neq 0, \theta = a) = \frac{1}{3}$.

- S likes higher actions; R takes $a \neq 0$ iff $\Pr(\theta = a) \geq \frac{3}{4}$.

Since S must get same payoff from all equilibrium messages and can't always make R willing to take action 2, S's equilibrium payoff can't exceed 1.

S can get 1 with $M = \{1, 2\}$, $\sigma(2|1) = \frac{1}{3}$, $\sigma(2|2) = 1$.

- Then R takes action 1 after message 1, R is indifferent after message 2 and mixes 50-50.

If S had commitment power, would be an eqm for S to commit to $\sigma(2|1) = \frac{1}{3}$, $\sigma(2|2) = 1$, R to take action 2 after message 2.

- S's commitment payoff is $\frac{1}{3} \cdot 1 + \frac{2}{3} \cdot 2 = \frac{5}{3}$.

Lipnowski Ravid: Main Result

For any posterior belief of the receiver $\tilde{\mu} \in \Delta(\Theta)$, let $v(\tilde{\mu}) = u_S(a_R(\tilde{\mu}))$, the sender's indirect utility from inducing this posterior. This defines a function $v : \Delta(\Theta) \to \mathbb{R}$.

Let $\bar{v} : \Delta(\Theta) \to \mathbb{R}$ and $\hat{v} : \Delta(\Theta) \to \mathbb{R}$ be the quasi-concave envelope and the concave envelope of v, respectively.

- $\bar{v}$ is the smallest quasi-concave function that is everywhere greater than v; similarly for $\hat{v}_S$.
- For all $\tilde{\mu}$, we have $v(\tilde{\mu}) \leq \bar{v}(\tilde{\mu}) \leq \hat{v}(\tilde{\mu})$.

Lipnowski-Ravid's main result is

Theorem

The sender's best equilibrium payoff is $\bar{v}(\tilde{\mu})$.

Comparison to Bayesian Persuasion

Kamenica Gentzkow 11 (cf. Aumann Maschler 95) show that if the sender can **commit** to any communication rule (conditional distribution over messages; experiment), her best equilibrium payoff is $\hat{v}(\mu)$.

- Can't do better than $\hat{v}(\mu)$, because $\mathbb{E}[\tilde{\mu}] = \mu$ by "Bayes' plausibility" (beliefs are martingale), and hence

$$\mathbb{E}[v(\tilde{\mu})] \underbrace{\leq}_{v \leq \hat{v}} \mathbb{E}[\hat{v}(\tilde{\mu})] \underbrace{\leq}_{\hat{v} \text{ concave, Jensen}} \hat{v}(\mu).$$

- Can get $\hat{v}(\mu)$, because there exists a distribution of $\tilde{\mu}$ such that $\mathbb{E}[\tilde{\mu}] = \mu$ and $\mathbb{E}[v(\tilde{\mu})] = \hat{v}(\mu)$, and any such distribution is induced by some experiment. (This is called the **splitting lemma**, usually attributed to Aumann-Maschler.)

Lipnowski Ravid (cntd.)

Intuition for Lipnowski Ravid's result:

- Can't do better than $\bar{v}(\mu)$, because this requires a communication rule that induces $\tilde{\mu}, \tilde{\mu}'$ s.t. $v(\tilde{\mu}) < v(\tilde{\mu}')$. Such a rule is not IC for Sender.
- Can get $\bar{v}(\mu)$, because there exists a distribution of $\tilde{\mu}$ such that $\mathbb{E}[\tilde{\mu}] = \mu$ and $v(\tilde{\mu}) = \bar{v}(\mu)$ for all $\tilde{\mu}$ in the support, and any such distribution is induced by some communication rule (and such a rule is IC for Sender).

Comparing to Bayesian persuasion, we see that the sender's "value of commitment" is $\hat{v}(\mu) - \bar{v}(\mu)$.

If no-disclosure is suboptimal with commitment, Sender typically does strictly better with commitment than without.

Purification

Critique of cheap talk with state-independent sender preferences: informative equilibria may not be purifiable.

Recall: a mixed equilibrium is **purifiable** if it is a limit of pure equilibria in games with independent payoff perturbations.

In CH, LR, sender mixes with different probabilities depending on θ, which she doesn't care about.

If perturb sender's preferences on A, it seems that her messages should be driven by these perturbations, not θ.

So robustness of informative equilibria in CH, LR is not clear.

Investigating this issue could be a good paper topic.

Multiple Senders

In many settings a decision-maker (Receiver) gets information from multiple advisors (Senders).

- Does “competition” between the senders leads to more info transmission?
- If so, limited info transmission results of CS and others may be questioned.

This question is studied by Battaglini 02, Ambrus Takahashi 08.

Shooting the Senders

Say an action a is **rationalizable** if $a = a_R(\tilde{\mu})$ for some $\tilde{\mu} \in \Delta(\Theta)$.

Suppose that there is a rationalizable action a^* such that, for each sender S_i ($i \in \{1, 2\}$) and each state θ, $u_{S_i}(a^*, \theta) \leq u_{S_i}(a, \theta)$ for every rationalizable action a.

- Not necessarily a realistic assumption, but illustrative.

Then there always exists a fully-revealing PBE:

- Senders announce θ honestly.
- If they agree, Receiver takes $a = a_R(\theta)$.
- If they disagree, Receiver believes that $\theta \sim \tilde{\mu}$ that rationalizes a^*, and takes a^*.

In the CS model with 2 senders with different biases, there is no single "shoot the senders" action that works, but a similar argument works whenever Θ and A are large compared to the senders' biases: if $m_1 \neq m_2$, Receiver believes that θ is extreme and takes an extreme action that is bad for both senders.

Ambrus Takahashi 08

Ambrus Takahashi extend this logic to give general conditions for the existence of a fully-revealing equilibrium.

Theorem

There exists a fully-revealing PBE iff, for all pairs of states $\theta_1, \theta_2 \in \Theta$, there exists a rationalizable action $a \in A$ such that $u_i \left(a_R \left(\theta_j\right), \theta_j\right) \geq u_i \left(a, \theta_j\right)$ for all $i \neq j \in \{1, 2\}$.

- Suppose a is taken when the reports are (θ_1, θ_2).
- The condition says that when the true state is θ_j, player i would rather report it truthfully and get $a_R \left(\theta_j\right)$ than misreport it as θ_i and get a.

Robust Equilibria

"Shoot the senders" equilibria do not seem very robust.

Battaglini considers a refinement where each Sender observes a random state $\tilde{\theta}$ (e.g., $\tilde{\theta} \sim \textit{Uniform}\,(\Theta)$) with prob ε and takes $\varepsilon \rightarrow 0$.

If the senders' biases are sufficiently large, there is no robust fully-revealing PBE in the 1-dimensional CS model.

Battaglini's main result: for generic Sender biases, there is a robust fully-revealing PBE in the d-dimensional CS model for any $d \geq 2$.

Full Revalation for d>=2

Suppose $\Theta = A = \mathbb{R}^d$, $u_R(a,\theta) = -\sum_{j=1}^{d}(a_j - \theta_j)^2$, $u_{S_i}(a,\theta) = -\sum_{j=1}^{d}(a_j + b_{i,j} - \theta_j)^2$, where $b_i \in \mathbb{R}^d$ is Sender i's bliss point.

Theorem

For any $d \geq 2$ and any b_1, b_2 such that $b_1 \neq \alpha b_2$ $\forall \alpha \in \mathbb{R}$, there exists a robust fully-revealing PBE.

- Each Sender i reports projection of θ onto the hyperplane orthogonal to b_i.
- When b_1, b_2 are linearly independent, the reports pin down θ, and Receiver takes $a_R(\theta)$.
- Each Sender is indifferent among all reports, and so is willing to report truthfully.
- Also, every pair of reports is sent on-path, so the equillibrium is robust.

Another Result from Ambrus Takahashi 08

Ambrus Takahashi's motivation was that Battaglini's argument relies on Θ and A being unbounded (or bounded product sets).

- Suppose $\Theta = A = \{\theta \in \mathbb{R}^2 : \theta_1 + \theta_2 \leq 1$. Not all combinations of $\hat{\theta}_1, \hat{\theta}_2$ are possible.
- Suppose $b_1 = (1, 0)$, $b_2 = (0, 1)$. Intuitively, if we try to do something similar to Battaglini while respecting feasibility, S_1 will shade down report of θ_2 to make larger θ_1's feasible.

They show that if $A = \Theta$ is a smooth compact set in $\mathbb{R}^d$, $d \geq 2$, $a_R(\theta) = \theta$, and b_1 and b_2 are not co-linear, then there is no robust fully-revealing PBE.

Remark: Battaglini and Ambrus-Takahashi both focus on when a (robust) fully-revealing PBE exists. This is an extreme type of result. More realistically, there should only be partial info transmission even with multiple senders. There seems to be little work on this.

Multiple Receivers

Cheap talk with multiple receivers also raises new issues. For example, should communication be public or private?

- Public communication can improve credibility. If Sender is biased up compared to Receiver 1 and down compared to Receiver 2, truthful communication may be IC when it's public.
- But private communication is more flexible. If Sender and Receiver 1 have the same preferences but Receiver 2 is very biased, with private communication can have info transmission to Receiver 1 together with babbling to Receiver 2, but with public communication the only equilibrium may be babbling.

These issues are studied by Farrell Gibbons 89, Goltsmann Pavlov 11. We skip these papers.

Multiple Rounds of Communication

The CS model restricts to 1 round of communication.

- At first this appears to be without loss, since Receiver doesn't have private info.
- However, random messages from Receiver can serve as a correlating/randomization device. Can this help?
- If so, can it be even better to have many rounds of communication and randomization? How would this work?

This is studied by Krishna Morgan 04, Aumann Hart 03.

Jointly Controlled Lotteries

An important idea in both papers is that of a jointly controlled lottery (which goes back to Aumann Maschler in the 60s).

By simultaneously sending cheap talk messages, players can simulate a public randomizing device in an incentive compatible manner.

- Suppose each player i names a number $n_i \in \{0, \dots, 99\}$, and we define the "output" of the pair (n_1, n_2) as $n_1 + n_2 \bmod 100$.
- Suppose each player is expected to choose n_i uniformly at random. Then, from the perspective of each player i, the output is uniform on $\{0, \dots, 99\}$ regardless of her choice of n_i.
- So, regardless of how the output maps into the players' future play, it is optimal for each player to choose $n_i \sim Uniform\{0, \dots, 99\}$ when she believes the other is doing so as well.
- The players can therefore use cheap talk to simulate public randomization.

Krishna Morgan 04

Krishna Morgan show that in the CS linear-quadratic model, efficiency can be improved if first Receiver sends a message, then public randomization occurs, and then Receiver sends another message.

- Using jointly controlled lotteries, can replace public randomization with a round of "face-to-face communication" between the parties.
- At first glance this is quite surprising, as there is no "obvious" role for multi-round communication in CS.

Intuition: Subsequent public randomization exposes Sender to risk: she may get a chance to reveal more info, or she may not.

- With convex loss functions, this risk is especially painful for extreme Sender types.
- So extreme Sender types will "confess" to avoid exposure to risk. The remaining types can then communicate a la CS when public randomization allows them to do this.
- Sometimes on net this beats the best CS equilibrium.

Krishna Morgan: Example

Suppose $b = \frac{1}{10}$. The best CS equilibrium partitions $[0, 1]$ into two intervals, $\left[0, \frac{3}{10}\right]$ and $\left[\frac{3}{10}, 1\right]$.

Consider the following multi-round equilibrium:

- First, Sender reveals whether θ is greater or less than $\frac{1}{5}$.
- Then, public randomization occurs and either *succeeds* (prob p) or *fails* (prob $1 - p$).
- If it succeds and $\theta > \frac{1}{5}$, Sender further reveals whether θ is greater or less than $\frac{2}{5}$.

If $p = 1$, we would be back to CS, with partition $\left\{\left[0, \frac{1}{5}\right], \left[\frac{1}{5}, \frac{2}{5}\right], \left[\frac{2}{5}, 1\right]\right.$. This is not an equilibrium, because type $\frac{1}{5}$ Sender strictly prefers $\left[\frac{1}{5}, \frac{2}{5}\right]$.

But with $p < 1$, type $\frac{1}{5}$ Sender takes a risk if she reports $\theta > \frac{1}{5}$ in Round 1, because if pub rand fails she gets stuck with $\left[\frac{1}{5}, 1\right]$.

- Can calculate that setting $p = \frac{5}{9}$ makes the type $\frac{1}{5}$ Sender indifferent, and the resulting equilibrium Pareto-dominates the best CS equilibrium.

Remarks

Krishna Morgan show that an equilibrium of this form Pareto-dominates the best CS equilibrium whenever $b < \frac{1}{8}$.

When $b \geq \frac{1}{8}$, equilibria of this form still exist, but they are less efficient that the best CS equilibrium.

When $b \in \left(\frac{1}{8}, \frac{1}{\sqrt{8}}\right)$, Krishna Morgan construct a different, non-monotone equilibrium that Pareto-dominates the best CS equilibrium.

Aumann Hart 03

Aumann Hart consider general (finite) 2-player games where one player has private info.

- Allow unboundedly many rounds of simultaneous bilateral cheap talk, prior to play of a 1-shot game.
- Finiteness implies doesn't formally nest CS / Krishna Morgan, but conceptually much more generally.

Aumann Hart 03: Main Results

1. Without loss to consider strategies where in odd periods Sender unilaterally sends a message to Receiver, in even periods Sender and Receiver play a jointly controlled lottery.
    - Krishna Morgan's construction had 3 periods: message, JCL, message.
2. Geometric characterization of the set of PBE payoffs in terms of the "di-span" of the NE payoffs of the game without cheap talk.
    - This is a convexification-type procedure, somewhat similar to the quasi-convex envelope in Lipnowski Ravid. (Question: is there a simple way to explain the relationship?)
    - Much more general than Krishna Morgan's analysis, but the characterization seems hard to apply outside simple examples. So the papers are complementary.

Mediated Communication

Cheap talk is a fairly extreme model, in that there are no available commitment devices that might improve credibility. We can learn something about how different "communication institutions" work by comparing them with cheap talk.

Goltsman Hörner Pavlov Squintani 09 compare three communication protocols in the CS linear-quadratic setting:

- **Negotiation:** Long cheap talk, as in Krishna Morgan 04
- **Mediation:** There is a trusted mediator with commitment power. Sender sends message to mediator, who then sends message to Receiver. (Equivalent to designing a system to add noise to Sender's message.) This can relax Sender's IC.
- **Arbitration:** Like mediation but now the mediator directly takes a decision, rather than communicating to Receiver. This relaxes Receiver's IC in addition to Sender's.
    - This is equivalent to delegating decision to Sender.

Comparing Communication Protocols

In general, the set of implementable equilibrium outcomes (mappings $\Theta \to \Delta(A)$) under different communication protocols satisfies

$$
\begin{aligned}
& \text{no communication: } a \text{ independent of } \theta \\
\subset \; & \text{1-shot cheap talk} \\
\subset \; & \text{negotiation} \\
\subset \; & \text{mediation} \\
\subset \; & \underbrace{\left\{ \begin{array}{c} \text{arbitration/delegation,} \\ \text{Bayesian persuasion} \end{array} \right\}}_{\text{non-nested}} \\
\subset \; & \text{"integration": arbitrary } a(\theta).
\end{aligned}
$$

Arbitration/Delegation

An **arbitration rule** is a mapping $a_R : \Theta \to \Delta(A)$ satisfying Sender IC: $u_S(a_R(\theta), \theta) \geq u_S(a_R(\theta'), \theta)$ for all θ, θ'.

In the CS setting, the arbitration rule that maximizes Receiver's expected utility takes the form of **upward censorship**: $a_R(\theta) = \min\{\theta + b, \bar{a}\}$ for some $\bar{a} \in [0, 1]$.

- Sender can choose any action she pleases up to $\bar{a}$.
- The solution is intuitive: since Sender is upward-biased, Receiver should prohibit the highest actions (while prohibiting intermediate actions just adds variance).
- This was first proved by Holmström 77 and Melumad Shibano 91 (cf. Alonso Matouschek 08), albeit restricting to deterministic arbitration rules. GHPS show that in the uniform-quadratic model, it's optimal to set $\bar{a} = 1 - b$ if $b < \frac{1}{2}$, and otherwise set $\bar{a} = \frac{1}{2}$ (i.e., only allow action $\frac{1}{2}$; this is the outcome of the babbling equilibrium).
- Can be proved using the envelope characterization of Sender IC (e.g., Myerson 81).

Mediation

As compared to arbitration, mediation adds Receiver IC:
$\mathbb{E}\left[u_R\left(a,\theta\right)|a\right] \geq \mathbb{E}\left[u_R\left(a',\theta\right)|a\right]$ for all on-path a, all a'.

For $b < \frac{1}{2}$ (the non-trivial case), GHPS show that there exists an optimal mediation rule where the mediator randomizes between two actions in each state, with probability independent of the state: there exists a probability $p \in (0,1)$, intervals $\{[\theta_{n-1},\theta_n]\}_{n=1}^N$, and actions $\{a_n\}_{n=1}^N$ such that, when $\theta \in [\theta_{n-1},\theta_n]$, the mediator recommends action b with prob p and recommends action a_n with prob $1-p$, where $a_1 = b$ and $a_n > b\ \forall n > 1$.

Mediation: Remarks

- As in the 2-stage equilibrium of Krishna Morgan 04, introducing randomness relaxes Sender's IC constraint, at the cost of making the final message received by Receiver less informative.
- The same mechanism appeared earlier in Blume Board Kawamura 08, who interpreted it as introducing a probability p that Sender's message gets lost. They showed that such "noisy talk" can improve on the best CS equilibrium. But they didn't show that this is actually an optimal equilibrium with arbitrary mediation.
- The optimal mechanism of BBK and GHPS is not uniquely optimal. When $b < \frac{1}{8}$, Krishna-Morgan's 2-stage equilibrium gives the same expected payoffs.

Negotiation/Long Cheap Talk

GHPS show that Krishna-Morgan's 2-stage equilibrium gives the same payoffs as optimal mediation when $b < \frac{1}{8}$, so in this case optimal negotiation and optimal mediation coincide.

When $b > \frac{1}{8}$, they show that negotiation cannot attain the optimal mediation solution.

A characterization of the optimal negotiation payoff when $b > \frac{1}{8}$ remains unknown, as well as whether there is any non-babbling equilibrium when $b > \frac{1}{\sqrt{8}}$ (to my knowledge).

More References

Corrao Dai 2023 study mediation in general cheap talk games with state-independent sender preferences.

- Corrao 2023 studies mediation with transfers.

Whitmeyer 2024 (and others) studies optimal mediation from the Receiver's perspective.

Ivanov 2010 studies **information control**: optimal choice of Sender's info structure prior to cheap talk.

Kolotilin Li Li 2013 study **limited principal authority:** Receiver cannot delegate the action to Sender, but can commit in advance to restrict his own action set to some $\tilde{A} \subset A$.

Kolotilin Zapechelnyuk 2019 study connections between optimal delegation and Bayesian persuasion.