Signaling Games

Game Theory · Module 4 — Information And Beliefs

Node 4.5

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Information and Beliefs **Node:** 4.5

A signaling game is the canonical two-stage Bayesian game: a **sender** with private type chooses a **signal**, then a **receiver** observes only the signal (not the type) and chooses a **response**. The sender's type, the signal, and the receiver's response jointly determine both players' payoffs. Signaling games capture an enormous range of phenomena: Spence's education signaling of ability, costly advertising, warranty as quality signal, limit pricing, sovereign debt signaling. They are the first genuinely "dynamic" Bayesian games we encounter, solved via PBE (4.4) with belief refinements. In poker, every bet is a signal — you are the sender, your hand is the type, the sizing is the signal, and the opponent is the receiver updating beliefs about your range.

Formal Definition

A **signaling game** is a two-player game with the following structure:

1. **Nature** draws a type $\theta \in \Theta$ for the Sender according to common prior $p \in \Delta(\Theta)$.
2. **Sender** observes θ and chooses a message $m \in M$.
3. **Receiver** observes m (but not θ) and chooses an action $a \in A$.
4. **Payoffs**: $u_S(\theta, m, a)$ and $u_R(\theta, m, a)$.

Sender strategy. $s_S : \Theta \to \Delta(M)$, the distribution over signals by type.

Receiver strategy. $s_R : M \to \Delta(A)$, the response to each possible signal.

Belief system. $\mu : M \to \Delta(\Theta)$, the Receiver's belief about the Sender's type after observing each message.

Perfect Bayesian equilibrium of a signaling game. A triple (s_S, s_R, μ) such that:

1. **Sender optimality.** For every θ and every m in the support of $s_S(\theta)$:

$$m \in \arg\max_{m' \in M} \mathbb{E}_{a \sim s_R(m')}[u_S(\theta, m', a)].$$

2. **Receiver optimality.** For every m in the support of s_S (image of Sender's strategy):

$$s_R(m) \in \arg\max_{a \in A} \sum_{\theta} \mu(\theta \mid m) \cdot u_R(\theta, m, a).$$

3. **Bayesian consistency (on-path).** For every m reached with positive probability:

$$\mu(\theta \mid m) = \frac{p(\theta) \cdot s_S(\theta)(m)}{\sum_{\theta'} p(\theta') \cdot s_S(\theta')(m)}.$$

Classification of equilibria:

- **Separating:** distinct types send distinct messages; the Receiver learns the type exactly from the message.
- **Pooling:** all types send the same message; the Receiver learns nothing beyond the prior.
- **Semi-separating (partial pooling):** some types share messages, others separate.
- **Babbling:** messages are uninformative (equilibrium may exist with $\mu(m) = p$ for all m).

Intuition

A signaling game asks: when information is private and action is public, can the informed party credibly convey information through their actions?

The answer depends on the **single-crossing property** of payoffs: if different types have different marginal costs for sending the same message, there can be separating equilibria where each type picks a message that is cheaper for them than for other types. If types have identical cost structures, separation is impossible — every message is equally attractive across types — and only pooling or babbling survives.

Spence's education example is paradigmatic. High-productivity types can earn an education more cheaply (effort per year) than low-productivity types, so education is a "costly signal" that sorts types. In equilibrium, high types education themselves, low types do not, and the wage differential reflects the separating information.

A signaling game has *lots* of PBE because off-path beliefs are unconstrained. You can often support a pooling equilibrium by assuming "if the Sender deviated, the Receiver would think the Sender is the worst type," which deters any deviation. The **intuitive criterion** (Cho–Kreps, 1987) rules out unreasonable off-path beliefs by asking: "could this deviation benefit any type at all?" If the answer is "only type H could benefit from this deviation," then off-path belief should put probability 1 on H, not on L. Refinements shrink the equilibrium set dramatically.

For signaling to succeed (separate types), the signal must be **costly** and **differentially costly**: cheap signals can be faked by low types, so they don't separate. This is the "cheap talk" limit (Crawford–Sobel, 1982): when messages are costless, separation is much harder, and equilibria involve coarse partitioning of the type space.

Structural Role

Signaling games are the canonical two-stage Bayesian game and the simplest vehicle for PBE refinements:

- **4.1, 4.4 ⇒ 4.5.** Bayesian games + PBE → signaling games.
- **4.5 ⇒ 4.6.** Sequential equilibrium refines PBE via trembles; signaling games are where this matters most.
- **4.5 ⇒ 5.5 Chain store paradox.** Reputation is signaling across repeated interaction.
- **4.5 ⇒ 6.1 Mechanism design.** Screening (the dual of signaling) is the uninformed party designing the game; signaling is the informed party.
- **4.5 ⇒ 6.3 Auction theory.** Bid = signal of value; auction design is mechanism design for screening via bids.
- **4.5 ⇒ poker theory.** Every bet is a signal. Bet sizing conveys range information.

Mathematical Development

Separating equilibrium existence

A separating PBE in a two-type signaling game $\Theta = \{H, L\}$ with messages $M = \{m_1, m_2\}$ exists if there is an assignment $H \to m_H, L \to m_L$ with $m_H \neq m_L$ such that:

- Neither type prefers the other's message: $u_S(H, m_H, s_R(m_H)) \geq u_S(H, m_L, s_R(m_L))$ and similarly for L.
- The Receiver's optimal response given full revelation is $s_R(m_H)$ (believing $\theta = H$) and $s_R(m_L)$ (believing $\theta = L$).

This is the **single-crossing / sorting** condition: the message cost structure must differ between types in a way that makes the assignment incentive-compatible.

Pooling equilibrium existence

A pooling PBE has all types sending the same message m^*. Receiver belief on m^* is the prior: $\mu(\theta \mid m^*) = p(\theta)$. Receiver plays best response to prior.

For this to be PBE, off-path beliefs at other messages must deter deviation. Typically, $\mu(\theta \mid m') =$ some distribution that makes $s_R(m')$ sufficiently unattractive that no type wants to send m'.

The intuitive criterion

Definition. A PBE satisfies the intuitive criterion if for every off-path message m', the Receiver's belief $\mu(\theta \mid m')$ places probability zero on types θ for which

$$\max_{a \in A} u_S(\theta, m', a) < u_S(\theta, m^*, s_R(m^*))$$

where m^* is θ's equilibrium message. In words: if type θ can never benefit from sending m' (even under the most favorable Receiver response), then the Receiver should not think θ sent m'. Types that could conceivably benefit from a deviation are "candidates" for having made it.

The intuitive criterion rules out many pooling equilibria: if only the high type could benefit from an off-path education deviation, the off-path belief must concentrate on H, so the Receiver pays the H wage, so H deviates — the pooling breaks.

Cheap talk (Crawford–Sobel, 1982)

If the message is costless — $u_S(\theta, m, a)$ is independent of m — then every separating equilibrium requires that the Sender types strictly prefer their "assigned" message given the Receiver's response. Without cost differences, this is impossible for generic games: all types have the same payoff from every message-response combination, so any separating profile is indifferent.

Crawford–Sobel theorem. In a cheap-talk game where the Sender's private type is a real number $\theta \in [0, 1]$, the Sender sends a message, the Receiver chooses an action, and the Sender's bias is b (preferring higher actions by b), equilibria are finite "interval partitions": the Sender's type space is partitioned into intervals, and each interval sends the same message. The number of intervals is bounded by $1/(2b)$; for $b \geq 1/2$, only babbling equilibria exist.

The Crawford–Sobel result shows that cheap talk can convey *some* information but never full separation, and the amount of information that can be conveyed shrinks as interests diverge.

Other applications

- **Spence (1973).** Education as productivity signaling.
- **Milgrom–Roberts (1986).** Costly advertising as quality signaling (for "bad" products, advertising is not worth it at any response; "good" products advertise).
- **Kreps–Wilson, Milgrom–Roberts (1982).** Limit pricing as deterrence: incumbent monopolist signals low cost by pricing low, deterring entry.
- **Grossman (1981), Grossman–Hart (1981).** Full disclosure: under some conditions, the Sender reveals everything voluntarily because silence is interpreted as worst-case.

Worked Example

A two-type signaling game with both separating and pooling equilibria

$\Theta = \{H, L\}, p(H) = 0.5$. Messages $M = \{E, N\}$ (E = education, N = no education). Receiver chooses wage $w \in \{w_H, w_L\}$ with $w_H > w_L$. Sender payoff: $w - c(\theta, m)$ where $c(\theta, E)$ = education cost, $c(\theta, N) = 0$. Assume $c(H, E) = 2, c(L, E) = 8$. Receiver payoff: $-(w - \text{productivity}(\theta))^2$, productivity(H) = 10, productivity(L) = 4.

Separating equilibrium attempt. $s_S(H) = E, s_S(L) = N$. Beliefs $\mu(H \mid E) = 1, \mu(H \mid N) = 0$. Receiver best response: $s_R(E) = 10, s_R(N) = 4$.

Check Sender optimality: - H: E gives $10 - 2 = 8$; N gives 4. H prefers E. ✓ - L: E gives $10 - 8 = 2$; N gives 4. L prefers N. ✓

Separating PBE exists.

Pooling equilibrium attempt. Both types send N. Belief $\mu(H \mid N) = 0.5$. Wage $s_R(N) = 0.5{\cdot}10 + 0.5{\cdot}4 = 7$. Off-path E: if belief puts probability 1 on L, wage is 4, sender payoff H: $4 - 2 = 2 < 7$, L: $4 - 8 = -4 < 7$. Neither deviates. Pooling PBE exists.

Apply intuitive criterion. For the pooling on N, consider the off-path message E. Can L benefit from E under any Receiver response? Max $u_S(L, E, a) = 10 - 8 = 2 < 7$. L can never benefit. So the belief after E should put probability 0 on L, probability 1 on H. Then Receiver pays $w_H = 10$, and H's deviation payoff is $10 - 2 = 8 > 7$ — H deviates. Pooling PBE fails the intuitive criterion.

Conclusion. Only the separating PBE survives the intuitive criterion. This is the "Spence lesson": with the single-crossing condition, refinements select the separating equilibrium.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Sender's type θ	Your hole cards
Message m	Your bet sizing and action (bet small, bet big, check, overbet, shove)
Receiver action a	Opponent's response: call, fold, raise, by what amount
Sender strategy s_S	Your bet sizing strategy: what you do with each hand
Receiver strategy s_R	Opponent's response strategy: what they do vs each of your sizings
Belief μ	Opponent's range inference from your sizing — "if you bet big, what do you have?"
Separating equilibrium	Different sizings correspond to different ranges (polarized vs merged)
Pooling equilibrium	One-size strategy where all your hands bet the same amount — opponent learns nothing from sizing
Intuitive criterion	Solver refinement: beliefs off-path make sense given "what hands would actually want to do this?"

Concrete situation

Bet sizing as signal in a river spot. You are BTN, facing BB check on the river. Your options: bet 33% pot, bet 75% pot, overbet 150% pot, or check. Each sizing is a potential message.

Pooling strategy (one size). If you always bet 75% regardless of hand, your sizing conveys no information about your hand beyond "I'm betting." BB's range-vs-your-75% is your entire betting range. This is a pooling equilibrium: all types send the same message.

Separating strategy (polarized sizing). Solver outputs often split sizings: bet small (33%) with medium-value hands, overbet (150%) with nuts/air polarized. This is a separating (actually partial-separating) equilibrium: different types send different messages, and BB's range inference differs by sizing. On seeing 33%, BB's belief is "medium value"; on seeing 150%, BB's belief is "nuts or air."

Why separation works. The payoffs must satisfy a single-crossing condition. The solver's single-crossing is the EV structure: for nuts, overbetting is cheaper (BB folds more, so you gain by applying max pressure) than betting small. For medium value, overbetting is expensive (you get called only by better, losing more). So different sizings are "cheaper" for different hand categories — the essential signaling condition.

Mixed-strategy partial separation. In practice, most solver outputs are *partial* separation: at each sizing, some range of hands is mixing. This is because perfectly separating would be detectable and counterable; the mixing creates sufficient range ambiguity that BB cannot respond deterministically. It's a semi-separating equilibrium.

Receiver best response. BB, on seeing each sizing, computes the posterior over your hand and plays best response. This is literally the PBE receiver-optimality condition. A solver enforces it at every sizing. The output: "vs your 33%, BB folds 40%, calls 60%; vs your 150%, BB folds 65%, calls 30%, raises 5%." These frequencies are optimal given the belief, which is Bayes-consistent with your sizing strategy.

What this understanding buys you

Every bet size is a message in a formal signaling game. When you choose 33% vs 75% vs overbet, you are choosing a signal. The opponent is receiving the signal and updating their belief about your range. This is not metaphor — it is the formal structure of a signaling game.

Polarization is signaling theory at work. "Polarized range" means "this sizing carries a message: 'I have nuts or nothing.'" The solver polarizes because the separation earns EV: nut hands get called less (by putting opponent in a tough spot) and air hands get fewer folds (by bluffing credibly). Both effects rely on the opponent's Bayesian inference from sizing to range.

Balance = equilibrium condition, not an aesthetic preference. "Balanced" ranges are PBE strategies: the bluff-to-value ratio at each sizing must be such that the opponent is indifferent between calling and folding. This is the receiver-optimality condition of the signaling game. If your sizing is "unbalanced," the opponent exploits by always calling (if you're too value-heavy) or always folding (if you're too bluff-heavy).

Intuitive criterion explains why certain deviations don't work. If you try to bluff with a hand that would never benefit from the bluffing sizing (e.g., medium pair overbetting — medium pair doesn't want to fold out opponent's bluff-catchers and doesn't get called by worse), the intuitive criterion says opponent should put belief 0 on that hand for that sizing. So the opponent's response deters the deviation. Solvers implement this automatically: hands only mix into sizings where they benefit.

Cheap talk doesn't work at the poker table. Poker is not cheap talk — every action has payoff consequences. This is why bet sizing *can* carry information: sending a high signal costs real chips. In pure cheap talk (just chatting), you couldn't credibly convey hand strength via language alone (beyond very coarse partitions) because there's no cost structure separating types. But the moment you back the message with money, signaling becomes possible.

Cross-Links

- **4.1 Bayesian games** — base framework.
- **4.2 Types and type spaces** — Sender's private information.
- **4.4 Perfect Bayesian equilibrium** — solution concept.

- **4.6 Sequential equilibrium** — refinement via trembles.
- **5.5 Chain store paradox** — reputation as multi-stage signaling.
- **6.1 Mechanism design** — screening (dual of signaling).
- **6.3 Auction theory** — bid as signal of value.

Structural Compression

Invariant: A signaling game is a two-stage Bayesian game in which an informed Sender sends a message and an uninformed Receiver responds, with the message partially revealing the Sender's type through Bayes-consistent beliefs; equilibrium outcomes depend on whether single-crossing conditions allow costly signaling to separate types.

Mechanism: Single-crossing cost structure enables incentive-compatible separating equilibria; pooling and babbling equilibria arise when signaling is cheap or types are too similar; refinements (intuitive criterion, divinity) use off-path belief logic to select among multiple equilibria.

Cross-System Echo: Signaling games formalize the general phenomenon of "credible revelation through costly action" — seen in evolutionary biology (handicap principle: peacock tail is a costly signal of fitness), labor economics (education as productivity signal), advertising (brand investment as quality signal), and political communication (policy commitments as type signals). In each case, only actions that are differentially costly across types can separate types credibly.

Exercises

1. Solve a two-type signaling game with messages $\{E, N\}$ and your choice of single-crossing cost. Characterize all separating and pooling PBE.
2. Show that the intuitive criterion rules out pooling equilibria in the Spence job market model when productivity gap exceeds the high type's cost.
3. In a cheap-talk game with Sender bias $b = 0.1$, apply the Crawford–Sobel analysis to determine the maximum number of intervals and thus the coarsest separation.
4. Construct a signaling game with three types and find a semi-separating equilibrium in which two types pool on one message and the third separates.
5. Show that in the Milgrom–Roberts advertising model, high-quality firms signal through costly advertising that would be unprofitable for low-quality firms.
6. Model a poker river bet-sizing decision as a signaling game with three hand types (value, medium, bluff) and two sizings (small, big). Solve for a semi-separating PBE and interpret.
7. **Argue, structurally,** why single-crossing is the key condition for separating equilibria in signaling games. What goes wrong if the single-crossing property fails?

Game Theory · Module 4 — Information And Beliefs · Node 4.5
Concepts: bayesian-updating, incomplete-information, bayesian-nash-equilibrium
Prerequisites: 4.1, 4.4
Enables: 4.6, 5.5

hbar.university — own.your.cognition