Lecture 3: Signaling Games

Alexander Wolitzky

MIT

14.126, Spring 2024

Communication Games

Typical structure of a communication game: a **sender** with private information takes an action (or otherwise "sends a message") that is observed by a **receiver**, who then takes an action that affects both players' payoffs.

We cover two classes of communication games:

1. **Signaling games:** sender's action set is *exogenous* and *payoff-relevant* ("go to college or not," "have beer or quiche for breakfast").
2. **Cheap talk games:** sender's "action" is a *payoff-irrelevant* message ("cheap talk"), we ask what can happen for *any message set*.

Related Models (Later in the Course)

1. **Reputation formation** and **bargaining**: can view as kinds of long-run signaling games (typically, infinite horizon+patient players).
2. **Social learning:** typically studies outcome of long-run observational learning/"communication"; often (not always) assume myopic agents so communication is non-strategic.

Signaling Games

- Sender has type $\theta \in \Theta$ (private info), with full support prior $\mu \in \Delta(\Theta)$.
- Sender has exogenous finite signal (action) space S, receiver has finite action space A.
- Sender observes θ, then takes $s \in S$.
- Receiver observes s (but not θ), then takes $a \in A$.
- θ, s, and a are all payoff-relevant: utilities are $u_1 : \Theta \times S \times A \to \mathbb{R}$, $u_2 : \Theta \times S \times A \to \mathbb{R}$.

Strategies

- $\Sigma_1 = (\Delta(S))^{\Theta}$ is the set of sender behavior strategies.
- $\sigma_1(\cdot|\theta)$ is type θ's probability distribution over signals s.
- $\Sigma_2 = (\Delta(A))^{S}$ is the set of receiver behavior strategies.
- $\sigma_2(\cdot|s)$ is receiver's probability distribution over actions a after observing signal s.
- Denote a strategy profile by $\sigma = (\sigma_1, \sigma_2)$.

Best Response Correspondences

- $BR(p, s) = \text{argmax}_{a \in A} \mathbb{E}_p [u_2(\theta, s, a)]$ is set of receiver's BR's to s given *belief* $p \in \Delta(\Theta)$.
- $BR(\tilde{\Theta}, s) = \bigcup_{p \in \Delta(\tilde{\Theta})} BR(p, s)$ is set of receiver's BR's to s for *some* belief with support contained in $\tilde{\Theta} \subset \Theta$.
- $MBR(p, s)$ and $MBR(\tilde{\Theta}, s)$ are the corresponding sets for mixed best responses.

Some Applications of Signaling Games

Job market signaling: Player 1=student/worker, Player 2=the labor market, θ=worker's ability, s=worker's education, a=wage.

- Implication: workers get too educated to signal high ability.

Initial public offerings: Player 1=owner of private firm, Player 2=potential investors, θ=value of firm, s=fraction of company retained, a=price.

- Implication: firms retain too much stock to signal high value.

Monetary policy: Player 1=the Fed, Player 2=the market, θ=how much Fed cares about inflation, s=first period inflation, a=output or inflation expectations.

- Implication: Fed is too tough on inflation today to signal it will be tough in the future.

Perfect Bayesian Equilibrium

Defining PBE in general extensive-form games can be tricky, but there is a simple, standard definition of PBE in signaling games.

Definition

Strategy profile σ is a **PBE** if

1. For each $\theta \in \Theta$, $u_1(\theta, \sigma) = \max_{s \in S} u_1(\theta, s, \sigma_2(\cdot|s))$.
2. For each on-path signal s (i.e., $s \in \bigcup_{\theta \in \Theta} \operatorname{supp} \sigma_1(\cdot|\theta)$), $\sigma_2(\cdot|s) \in MBR(p(\cdot|s), s)$, where $p(\cdot|s)$ is the posterior belief given s obtained by Bayes' rule.
3. For each off-path signal s, $\sigma_2(\cdot|s) \in MBR(\Theta, s)$.

- (1) and (2): σ is a NE.
- (3): R's play at off-path info sets is sequentially rational for *some* belief.
- R's off-path beliefs aren't pinned down by Bayes' rule. For this reason, signaling games often have many equilibria.
- Refinements needed to restrict off-path beliefs.

Pooling, Separating, Semi-Separating

PBE in signaling games can be classified as pooling, separating, or semi-separating.

A PBE is **pooling** if all sender types take the same pure action s. Then receiver BR given s is $\sigma_2(\cdot|s) \in MBR(\mu, s)$.

- (Or maybe all types take the same mixed action.)

A PBE is **separating** if each sender type θ takes a different pure action $s(\theta)$. Then receiver BR given $s(\theta)$ is $\sigma_2(\cdot|s(\theta)) \in MBR(\delta_\theta, s(\theta))$.

- (Or maybe all types take mixed actions with disjoint supports.)

A PBE is **semi-separating** if it is not pooling or separating.

Job-Market Signaling (Spence 74)

- Player 1 is a worker, observes her productivity $\theta \in \mathbb{R}_+$, chooses education level $s \in [0, \infty)$.
- Finite set of possible types $\{\theta_1, \theta_2, \ldots, \theta_n\}$, with $0 < \theta_1 < \theta_2 < \ldots < \theta_n$.
- Player 2 is "the labor market," sees s but not θ, chooses a wage a.
- Player 2's payoff is $-(a-\theta)^2$. Shorthand for labor market setting wage equal to expected productivity, $a = \mathbb{E}[\theta|s]$.
- Player 1's payoff is $a - s/\theta$. Education is costly, but less so for more productive types.

PBE in the Job-Market Signaling Game

First consider separating PBE.

In a separating PBE, the lowest type θ_1 is paid θ_1.

- Since θ_1 is the lowest possible value of $\mathbb{E}[\theta|s]$ for any belief, type θ_1 must take $s = 0$ in any separating PBE.
- Given this, type θ_2 must take s_2 *at least* $(\theta_2 - \theta_1)\theta_1$. Otherwise, type θ_1 would rather take s_2 and get paid θ_2 than take 0 and get paid θ_1.
- Recursively, type θ_k must take $s_k \geq s_{k-1} + (\theta_k - \theta_{k-1})\theta_k$.

The strategy profile where these inequalities all hold with equality is the **least-cost separating equilibrium** (Riley 79). It's defined by specifying that the lowest type takes the action that maximizes their utility given that their type is revealed; the next-lowest type takes the action that maximizes their utility given that their type is revealed and the lowest type does not want to copy them; etc.

- This is the Pareto optimal separating eqm for the sender types.

PBE in the Job-Market Signaling Game (cntd.)

There is a continuum of other separating equilibria.

- Suppose player 2 believes that $\theta = \theta_1$ if $s < (\theta_2 - \theta_1)\,\theta_1 + \varepsilon$.
- Then it's optimal for type θ_2 to take $s = (\theta_2 - \theta_1)\,\theta_1 + \varepsilon$, because (for small ε) $\theta_2 > \theta_1$ implies that

$$\theta_2 - \left((\theta_2 - \theta_1)\,\theta_1 + \varepsilon\right)/\theta_2 > \theta_1.$$

- These equilibria are Pareto dominated by the least-cost separating eqm.

There is also a continuum of pooling equilibria.

- Suppose player 2 believes that $\theta = \theta_1$ if $s < \varepsilon$, otherwise $\theta \sim \mu$.

There are also semi-separating equilibria.

Clearly, need some refinements/selection to make sharp predictions.

Equilibrium Refinements in Signaling Games

- Banks Sobel 1987 and Cho Kreps 1987 were the first systematic attempts to justify particular equilibrium selection in signaling gmaes.
- Both papers drew inspiration from Kohlberg Mertens 1986 and motivated their refinements as consequences of KM's strategic stability.
- Cho Kreps also gave an informal motivation in terms of (unmodeled) "speeches" that deviating players could make to suggest how their deviations should be intepreted.
- "Stiglitz critique": if players can make such speeches, the speeches and the resulting inferences should be modeled as part of the game.

Responses to Equilibrium Multiplicity/Stiglitz Critique

We will cover the classic eqm refinements of Banks-Sobel and Cho-Kreps, which remain important even though they are subject to the Stiglitz critique and other critiques of forward induction-type reasoning.

There are also other approaches to equilibrium selection in signaling game that are not subject to the same critiques. We will cover some; more on this in 127.

1. Assume signals are observed with noise, so there are no off-path information sets. (Matthews Mirman 1983, Carlsson Dasgupta 1997)
2. Assume it is costly to observe signals, so off-path "belief threats" cannot support unreasonable pooling equilibria. (Denti 2021)
3. Derive signaling game refinements from models of non-equilibrium learning, where the meaning of off-equilibrium path signals is determined by the learning process (Fudenberg He 2018, Clark Fudenberg 2021).

The Intuitive Criterion (Cho Kreps 1987)

Fix a PBE σ and let $u_1(\theta,\sigma)$ be type θ's eqm payoff.

The set of sender types for whom a signal s is **not equilibrium dominated** is

$$\Theta^{IC}(s,\sigma) = \left\{\theta \in \Theta : \max_{a \in BR(\Theta,s)} u_1(\theta,s,a) \geq u_1(\theta,\sigma)\right\}.$$

- Set of types that can plausibly hope to do better by taking s. (Where "plausible"=receiver takes a BR to *some* belief.)

In *any* PBE, for *on-path* signals s, receiver's posterior $p(\cdot|s)$ must put probability 1 on $\Theta^{IC}(s,\sigma)$.

- A type $\theta \notin \Theta^{IC}(s,\sigma)$ cannot take s in any PBE, as this would give strictly less than $u_1(\theta,\sigma)$.

Idea of Intuitive Criterion: receiver should also put probability 1 on $\Theta^{IC}(s,\sigma)$ for off-path signals s.

Implicit "speech" behind this idea: "I'm sending s and you should believe that my type is in $\Theta^{IC}(s,\sigma)$, because otherwise I would have no reason to take s and make this speech."

Intuitive Criterion (cntd.)

Definition

A PBE σ *passes the Intuitive Criterion* if, for every $s \in S$ and $\theta \in \Theta$,

$$\min_{a \in BR\left(\Theta^{IC}(s,\sigma),s\right)} u_1\left(\theta, s, a\right) \leq u_1\left(\theta, \sigma\right).$$

- LHS is the minimum that type θ can expect from sending s if off-path $p\left(\cdot|s\right)$ puts prob 1 on $\Theta^{IC}\left(s, \sigma\right)$.
- Intuitive criterion says that type θ should expect at least this much from taking s, so this can't exceed θ's eqm payoff.

Intuitive Criterion in Job Market Signaling

Cho-Kreps show that the Intuitive Criterion selects the least-cost separating PBE in the Spence game if there are only 2 types.

First consider separating PBE.

- Signals $s > (\theta_2 - \theta_1)\,\theta_1$ are equilibrium dominated for type θ_1. So the IC requires that these signals are attributed to type θ_2.
- Therefore, type θ_2 never takes $s > (\theta_2 - \theta_1)\,\theta_1$.
- This rules out separating PBE with "excess signaling."

Intuitive Criterion in Job Market Signaling (cntd.)

Now consider PBE where both types send some $\hat{s}$ w/ prob >0.

- Note that receiver's BR is never more than θ_2.
- Let $\bar{s}$ be s.t.

$$\theta_2 - \bar{s}/\theta_1 = u_1(\theta, \sigma) = \mathbb{E}[\theta|\hat{s}] - \hat{s}/\theta_1.$$

  Signals above $\bar{s}$ are equilibrium dominated for type θ_1.
- Since type θ_1 is indifferent, type θ_2 strictly prefers wage θ_2 and signal $\bar{s} + \varepsilon$ to $u_2(\theta, \sigma)$, for small ε.
- Therefore, the eqm fails the Intuitive Criterion.

This argument fails with 3 or more types.

- For $\bar{s}$ to be equilibrium dominated for type θ_1 it must be unprofitable for type θ_1 even if it comes with wage θ_3.
- But if type θ_2 deviates to $\bar{s}$ she can only count on a wage above θ_2 (not θ_3). This may not exceed her eqm payoff.

More Restrictive Refinements (Banks Sobel 1987)

The set of mixed best responses to s that make type θ strictly prefer s to her eqm outcome is

$$D_\theta(s,\sigma) = \{\alpha \in MBR(\Theta, s) : u_1(\theta, s, \alpha) > u_1(\theta, \sigma)\}.$$

The set of mixed best responses that make type θ indifferent between s and her eqm outcome is

$$D^0_\theta(s,\sigma) = \{\alpha \in MBR(\Theta, s) : u_1(\theta, s, \alpha) = u_1(\theta, \sigma)\}.$$

Banks-Sobel propose that the receiver should put prob 0 on type θ after signal s if there is another type θ' such that every MBR to s that makes θ willing to deviate to s makes θ' strictly prefer to deviate to s: i.e., if $\exists\theta'$ s.t.

$$D_\theta(s,\sigma) \cup D^0_\theta(s,\sigma) \subseteq D_{\theta'}(s,\sigma).$$

Banks-Sobel 87 (cntd.)

Under this proposal, the set of types that the receiver can put prob >0 on is

$$\Theta^{D1}(s,\sigma) = \left\{ \theta : \forall \theta' \neq \theta,\ D_\theta(s,\sigma) \cup D^0_\theta(s,\sigma) \nsubseteq D_{\theta'}(s,\sigma) \right. .$$

Definition

A PBE σ *passes D1* if, for every $s \in S$, there exists $\alpha \in MBR\left(\Theta^{D1}(s,\sigma)\right)$ such that

$$u_1(\theta, s, \alpha) \leq u_1(\theta, \sigma) \text{ for all } \theta \in \Theta.$$

- Banks-Sobel also define a weaker variant of D1 called "divinity." This requires that deviations are deterred by beliefs that put less weight on types outside $\hat{\Theta}^{D1}(s,\sigma)$ than they get under the prior, rather than 0 weight.

D1 vs. Intuitive Criterion

D1 is more restrictive than the Intuitive Criterion in 2 ways.

1. D1 allows fewer receiver beliefs receiver after off-path s, so deviating sender can count on a higher payoff:
   If s is equilibrium dominated for θ, then
   $D_\theta(s,\sigma) = D_\theta^0(s,\sigma) = \emptyset$, so $\Theta^{D1}(s,\sigma) \subseteq \Theta^{IC}(s,\sigma)$.
   This makes D1 harder to pass.
2. For Intuitive Criterion, each deviating sender type posits the worst receiver BR for her own type.
   For D1, a single receiver BR must deter deviations by all sender types.
   This also makes D1 harder to pass.

D1 vs. Intuitive Criterion (cntd.)

Cho Kreps 1987 show that D1 selects the least-cost separating eqm in the Spence game for any number of types.

- If a signal s is sent with prob >0 by two types $\theta' > \theta$, then for any $s' > s$ we have $D_\theta(s', \sigma) \cup D^0_\theta(s', \sigma) \subseteq D_{\theta'}(s', \sigma)$, so type θ' breaks the PBE by sending s'.

Cho Sobel 1990 extend this result to all "monotonic" signaling games (all types have the same preferences over receiver MBR's, higher responses are better, single-crossing).

NWBR

A yet more restrictive condition is NWBR ("never a weak best response"; Kohlberg Mertens 1986, Cho Kreps 1987).

NWBR says the receiver should put prob 0 on type θ after signal s if every MBR that makes type θ indifferent to taking s makes *some* type θ' strictly prefer s: i.e., if

$$D^0_\theta(s,\sigma) \subseteq \bigcup_{\theta' \neq \theta} D_{\theta'}(s,\sigma)$$

Under NWBR, the set of types that the receiver can put prob >0 on is

$$\Theta^{NWBR}(s,\sigma) = \left\{ \theta : D^0_\theta(s,\sigma) \nsubseteq \bigcup_{\theta' \neq \theta} D_{\theta'}(s,\sigma) \right\}.$$

Definition

A PBE σ *passes NWBR* if, for every $s \in S$, there exists $\alpha \in MBR\left(\Theta^{NWBR}(s,\sigma)\right)$ such that

$$u_1(\theta, s, \alpha) \leq u_1(\theta, \sigma) \text{ for all } \theta \in \Theta.$$

NWBR vs. D1

Note that $\Theta^{NWBR}(s,\sigma) \subseteq \Theta^{D1}(s,\sigma)$, because under D1 a pair (θ, s) is deleted only if some type θ' prefers s for all MBR's, while under NWBR different types θ' could prefer s for different MBR's. Hence, NWBR is more restrictive than D1.

Results of Kohlberg-Mertens and Cho-Kreps imply that every signaling game has a PBE that satisfies NWBR.

Example: Hiring a Worker

θ_1	a_1	a_2	a_3
Hire	$16, 2$	$1, 0$	$-2, -1$
Pass	$0, 0$	$0, 0$	$0, 0$

θ_2	a_1	a_2	a_3
Hire	$8, 0$	$6, 1$	$-4, 0$
Pass	$0, 0$	$0, 0$	$0, 0$

θ_3	a_1	a_2	a_3
Hire	$4, -1$	$1, 0$	$-1, 1$
Pass	$0, 0$	$0, 0$	$0, 0$

- Sender is a firm, whose signal $s \in \{Hire, Pass\}$ is its choice of whether to hire a worker, who is the receiver.
- The firm's type is its quality, which can be high (θ_1), medium (θ_2) or low (θ_3).
- The worker's action is how hard she works: high (a_1), medium (a_2) or low (a_3).
- Worker wants to match effort to the firm's quality.
- All firm types have the same ordinal preferences $(Hire, a_1) \succ (Hire, a_2) \succ Pass \succ (Hire, a_3)$, but different cardinal preferences.

Hiring a Worker (cntd.)

θ_1	a_1	a_2	a_3
Hire	$16, 2$	$1, 0$	$-2, -1$
Pass	$0, 0$	$0, 0$	$0, 0$

θ_2	a_1	a_2	a_3
Hire	$8, 0$	$6, 1$	$-4, 0$
Pass	$0, 0$	$0, 0$	$0, 0$

θ_3	a_1	a_2	a_3
Hire	$4, -1$	$1, 0$	$-1, 1$
Pass	$0, 0$	$0, 0$	$0, 0$

There are PBE where every type plays *Pass*.

- Enforced by receiver playing a_3 when sender *Hires*, supported by belief that type θ_3 *Hires*.
- These PBE survive D1, because there is no single type that strictly prefers *Hire* whenever θ_3 does.
- But they do not survive NWBR, because whenever θ_3 weakly prefers *Hire*, *either* θ_1 or θ_2 strictly prefers *Hire*.

Mixed BR vs. Mixtures over Pure BRs

Given a set of types $\tilde{\Theta} \subset \Theta$, $\Delta\left(BR\left(\tilde{\Theta}, s\right)\right)$ is the set of mixtures over pure best responses to s and (possibly different) beliefs with support contained in $\tilde{\Theta}$. This is a larger set than $MBR\left(\tilde{\Theta}, s\right)$.

PBE and all refinements considered so far focus on $MBR\left(\tilde{\Theta}, s\right)$ for various $\tilde{\Theta}$'s, rather than $\Delta\left(BR\left(\tilde{\Theta}, s\right)\right)$.

Intuitively, this means that deviating senders know what the receiver believes following each off-path signal s.

This is part of the standard definition of PBE or sequential equilibrium, but it is not a necessary implication of equilibrium viewed as a stable outcome of a non-equilibrium learning model

- More on this in 127, where cover self-confirming equilibrium and related topics.

Mixed BR vs. Mixtures over Pure BRs (cntd.)

Fudenberg He 18 and Clark Fudenberg 21 derive new signaling game refinements as the outcomes of a non-equilibrium learning model with a large population of long-lived senders and receivers.

- This approach is not subject to the Stiglitz critique, because off-path signals are sent with positive probability during the learning process.

In Clark-Fudenberg, the learning model leads to a refinement called *justified communication equilibrium (JCE)*, which is the same as NWBR but with $MBR\left(\tilde{\Theta}, s\right)$ replaced by $\Delta BR\left(\tilde{\Theta}, s\right)$.

Clark-Fudenberg also show that JCE is more restrictive than the Intuitive Criterion in any signaling game, and that JCE, D1, and NWBR are all outcome-equivalent in *co-monotone* signaling games, where for each s all sender types have the same **cardinal** preferences over receiver actions.

Costly Monitoring (Denti 21)

Another recent approach to signaling game refinements assumes that the receiver must pay a (possibly small) cost to observe the signal.

So, now game is:

- Nature draws sender's type $\theta \in \Theta$, with full support prior $\mu \in \Delta(\Theta)$.
- Sender observes θ, then takes $s \in S$.
- Without observing θ or s, receiver chooses an *experiment* $P : S \to \Delta(X)$ from some set, where X is an exogenous outcome set.
- Receiver observes x distributed $P(\cdot|s)$, then takes $a \in A$.
- Utilities are $u_1(\theta, s, a)$ for sender, $u_2(\theta, s, a) - c(P)$ for receiver, where $c(\cdot)$ is a cost function defined on experiments.

No Free Information

To differentiate costly monitoring from the standard perfect monitoring model, assume **no free information**: if experiment P is feasible and experiment $Q \neq P$ is a Blackwell garbling of P, then Q is also feasible and $c(Q) < c(P)$.

This assumption is violated in standard signaling games, because the receiver can't save money by ignoring the signal.

No Free Information (cntd.)

No free information has the following important implication:

Lemma

If in any PBE the receiver chooses an experiment P such that $P(x|s) > 0$ for some $x \in X, s \in S$, then $P(x|s') > 0$ for some $s' \in S$ that is played with positive probability (i.e., $\exists \theta, s'$ s.t. $\sigma_1(s'|\theta) > 0$ and $P(x|s') > 0$).

- If signal x arises with positive probability only off-path, then the receiver can get the same joint distribution over $\Theta \times S \times A$ by deviating to an experiment Q that maps x to other signals. This experiment is a Blackwell garbling, so it is strictly cheaper.

Beer-Quiche Game (Cho-Kreps 87)

The beer-quiche game has two pooling equilibria: $(Quiche, Quiche, Don't)$ and $(Beer, Beer, Don't)$.

- $(Quiche, Quiche, Don't)$ fails the Intuitive Criterion, because *Beer* is equilibrium dominated for θ_w. So, IC selects $(Beer, Beer, Don't)$.

We will show that costly monitoring also selects $(Beer, Beer, Don't)$ when monitoring costs are sufficiently small (assuming no free information).

- In general, costly monitoring (with small monitoring costs) and the Intuitive Criterion don't always make the same predictions.

"Real Men Don't Eat Quiche"

Theorem

In any PBE of the costly monitoring game, for any cost function satisfying no free information, $\sigma_1 (Beer|\theta_s) = 1$.

Suppose that $\sigma_1 (Quiche|\theta_s) > 0$.

- θ_s eats *Quiche* only if $\Pr (Duel|Beer) > \Pr (Duel|Quiche)$. In particular, $\Pr (Duel|Beer) > 0$
- By the previous lemma, if $\Pr (Duel|Beer) > 0$ then $\Pr (Duel) > 0$. (Otherwise, receiver won't pay attention to signals that trigger *Duel*. This is the key difference from perfect monitoring.)
- ...

Proof (cntd.)

Next, by single-crossing, $\sigma_1 (Quiche|\theta_w) = 1$.

Hence, we have $\sigma_1 (Quiche|\theta_w) \geq \sigma_1 (Quiche|\theta_s)$ and $\Pr(Duel|Beer) > \Pr(Duel|Quiche)$. This implies that *Duel* is a signal of θ_w: i.e., by Bayes' rule,

$$\begin{aligned}
& \Pr(\theta_w|Duel) \\
= & \frac{\Pr(Duel|\theta_w)\Pr(\theta_w)}{\Pr(Duel|\theta_w)\Pr(\theta_w) + \Pr(Duel|\theta_s)\Pr(\theta_s)} \\
= & \frac{\Pr(Duel|Quiche)\Pr(\theta_w)}{\Pr(Duel|Quiche)\Pr(\theta_w) + \Pr(Duel|\theta_s)\Pr(\theta_s)} \\
\leq & \frac{\Pr(Duel|Quiche)\Pr(\theta_w)}{\Pr(Duel|Quiche)\Pr(\theta_w) + \Pr(Duel|Quiche)\Pr(\theta_w)} \\
= & \Pr(\theta_w) = 1/10.
\end{aligned}$$

But if $\Pr(\theta_w|Duel) < 1/2$, receiver should deviate to $Don't$.

PBE with Costly Monitoring

Since $\sigma_1\left(Beer|\theta_s\right) = 1$, there can be two kinds of PBE with costly monitoring:

1. Separating PBE where θ_s takes *Beer* and θ_w takes *Quiche*.
   - This PBE exists if monitoring costs are high, so that receiver chooses not to monitor and always takes $Don't$.
2. Semi-separating PBE where θ_s takes *Beer* and θ_w mixes.
   - This PBE exists if monitoring costs are lower. Receiver monitors s with just enough noise to make θ_w indifferent between *Beer* and *Quiche*; θ_w randomizes to give receiver incentive to acquire information.

There cannot be a pooling PBE on *Beer*, because receiver wouldn't monitor, and then θ_w would deviate to *Quiche*. However…

Vanishing Monitoring Costs

When monitoring costs vanish ($c(P) < \varepsilon$ for all P, including perfect monitoring, $\varepsilon \to 0$), all semi-separating PBE converge to $(Beer, Beer, Don't)$: that is, $\sigma_1 (Beer|\theta_w) \to 1$.

- Recall that θ_s always takes *Beer*, θ_w takes *Quiche* with prob that makes receiver just willing to monitor.
- When monitoring costs vanish, if $\sigma_1 (Quiche|\theta_w)$ does not also vanish, receiver strictly prefers to monitor.
- Hence, $\sigma_1 (Quiche|\theta_w)$ must also vanish.

Flexible Information Acquisition in General Games

A recent paper by Denti and Ravid (2023) study flexible information acquisition in general games.

- Payoff-relevant parameter θ, payoff-irrelevant parameter z (correlated with θ).
- Players simultaneously choose conditionally independent signals of (θ, z).
  - May not be conditionally independent given θ only. Indeed, the point of z is to allow correlated signals about θ.
- Can choose any signal, under cost function satisfying no free information.

Denti-Ravid show that the set of Bayes NE outcomes that arise for some cost functions is the set of **separated BCE**, where a BCE is **separated** if each player's best replies to distinct signals are disjoint.

- Similar logic as above: if the BR's corresponding to distinct signals overlap, strictly better to pool these signals and play the common BR.