543: Game Theory **Information Design**	**Week 7**

Updated September 24, 2021.

Preliminary and incomplete.

Thanks to 534 Fall 2020 students for proofreading! (Remaining errors are mine.)

Source(s): [Kamenica and Gentzkow, 2011], [Bergemann et al., 2015], [Roesler and Szentes, 2017].

1 Communication with Commitment

AKA Bayesian Persuasion. From [Kamenica and Gentzkow, 2011].

Example:

- There is a defendant, a prosecutor, and a judge
- States of the world: defendant is either guilty or innocent
- Judge actions: acquit or convict
- Judge utility = 1 if matches state, = 0 otherwise
- Prosecutor utility = 1 if convict, = 0 if acquit
- Shared prior belief $Pr[guilty] = 0.3$, $Pr[innocent] = 0.7$

Can the prosecutor convince the judge?

- If the prosecutor knows the state, then babbling

What if the prosecutor could choose a test before knowing the state?

- Prosecutor chooses an "investigation" π: $\pi(\cdot|guilty)$, $\pi(\cdot|innocent)$.

- If π uninformative: always acquit, prosecutor utility $= 0$
- If π fully informative: convict w.p. 0.3, prosecutor utility $= 0.3$
- Consider $\pi(i|innocent) = 4/7$, $\pi(g|innocent) = 3/7$, $\pi(i|guilty) = 0$, $\pi(g|guilty) = 1$.
- Judge posterior given i:
$$Pr[guilty|i] = Pr[i|guilty]Pr[guilty]/Pr[i] = 0$$
Acquit
- Judge posterior given g:
$$\begin{aligned} Pr[guilty|g] &= Pr[g|guilty]Pr[guilty]/Pr[g] \\ &= 1 \times 0.3/(1 \times 0.3 + 3/7 \times 0.7) = 0.5. \end{aligned}$$
Convict
- Prosecutor utility $= Pr[g] = 0.6$.

Is this the best investigation? Yes. We will show this below.

Each signal realization s induces posterior $q \in [0, 1]$

Each signal π induces distribution over posteriors.

In the examples above, with prior $p = 0.7$

- Uninformative: $q = 0.7$ with probability one
- Fully informative: $q = 0$ with probability 0.3, $q = 1$ with probability 0.7
- The optimal signal: $q = 0.5$ with probability 0.6, $q = 1$ with probability 0.4

Notice that in each case, $E[q] = p$.

Claim: For any signal π, $E[q] = p$. For any distribution μ over q such that $E_\mu[q] = p$, there exists a signal such that the distribution of posteriors is μ.

We say that μ is Bayes plausible if $E_\mu[q] = p$

So the problem is
$$\max_{\mu \text{ s.t. } E_\mu[q]=p} Pr[q \leq 0.5].$$

The objective can be written as $E_\mu[v(q)]$, where $v(q)$ is defined to be 1 if $q \leq 0.5$, and 0 if $q > 0.5$.

Let OPT be the value of the optimal solution. We will show that $OPT = 0.6$.

Define v' to be the lowest concave function that is pointwise at least as large as v. Let OPT' be the value of the optimal solution to the problem of maximizing $E_\mu[v'(q)]$ over all Bayes plausible posteriors μ.

Claim: $OPT \leq OPT'$.

- Proof: Since $v \leq v'$, we have $E_\mu[v(q)] \leq E_\mu[v'(q)]$ for all μ.

Claim: $OPT' = 0.6$

- Proof: Since v' is concave, in the optimal μ we have $q = p$ with probability 1.
- Therefore $OPT' = v'(p) = 0.6$.

Our analysis can be generalized to beyond the example:

- There is a sender and a receiver. The receiver is to make a decision.
- The payoffs depends on the decision as well as a random state of the world, with a shared prior p.
- Before knowing the state, the sender chooses a test.
- Upon observing the result of the test, the receiver updates her beliefs and makes a decision.

The solution is a generalization of before:

- For any posterior q, define $v(q)$ to be the sender's payoff from inducing a posterior belief q
- Define v' to be the "concavification" of v (lost concave function such that $v \leq v'$)
- The value of the optimal solution is $v'(p)$
- The optimal test induces a set of posteriors such that $v'(q) = v(q)$ and $E_\mu v(q) = v'(p)$.

2 The Limits of Price Discrimination

From [Bergemann et al., 2015].

Monopolist selling single good (zero cost) to a single consumer (or a population of consumers).

Example:

- $V = \{1, 2, 3\}$

- Uniform distribution: $f(1) = f(2) = f(3) = 1/3$
- optimal price is 2
- Producer surplus $PS = 4/3$
- Consumer surplus $CS = 1/3$

The surplus pair looks as follows:

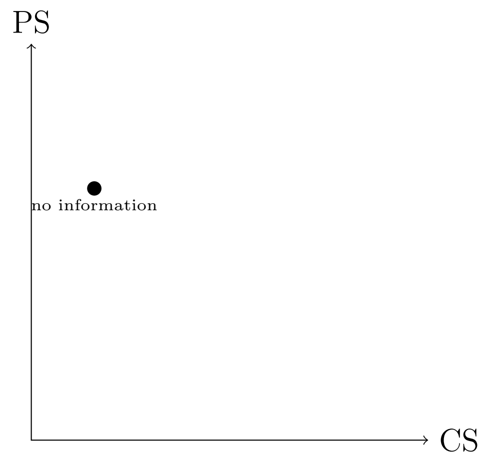

What if producer could get extra information?

- What if it could observe value perfectly?
- First degree price discrimination
- EG $PS = 2$, $CS = 0$.

The surplus pair looks as follows:

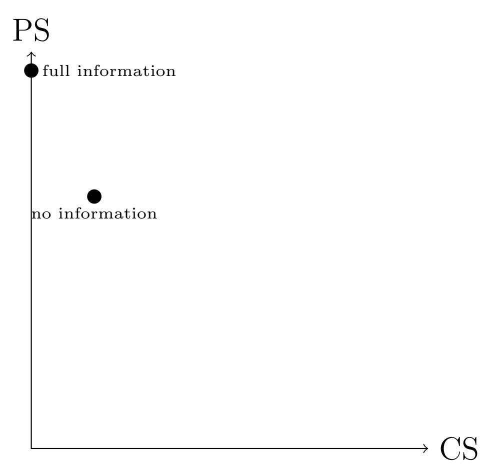

Notice that the expectation of the posteriors is equal to the prior:

$$1/3(1,0,0) + 1/3(0,1,0) + 1/3(0,0,1) = (1/3, 1/3, 1/3).$$

Claim: The distribution μ over posteriors must be Bayes plausible, that is, for any value v, $E_\mu[f'(v)] = f(v)$. Also any Bayes plausible distribution over posteriors is induced by some signal.

What is the set of all possible payoff pairs, when we consider all possible information that the seller can receive?

Three constraints:

- $CS \geq 0$
- $CS + PS \leq 2$
- $PS \geq 4/3$

Claim: The set of all possible payoff pairs = triangle defined by the above three inequalities

The set of all possible payoff pairs looks as follows:

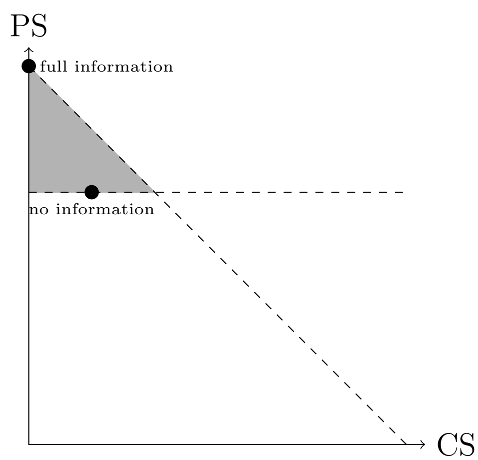

Bottom right corner is achieved with the following distribution over posteriors. For each posterior, the underlined prices are optimal for the seller.

- With probability 2/3: $f_1 = (\underline{1/2}, \underline{1/6}, \underline{1/3})$
- With probability 1/6: $f_2 = (0, \underline{1/3}, \underline{2/3})$
- With probability 1/6: $f_3 = (0, \underline{1}, 0)$

Verify that these posteriors average to the prior:

$$2/3(1/2, 1/6, 1/3) + 1/6(0, 1/3, 2/3) + 1/6(0, 1, 0)$$
$$= (1/3, 1/3, 1/3).$$

Verify that the profit = 4/3:

- f_1: revenue = 1
- f_2: revenue = 2
- f_3: revenue = 2
- Expected revenue = 4/3
- Other way to see: price = 2 is optimal for each distribution. So the seller doesn't benefit from information.

Consumer surplus:

- f_1: price = 1, $CS = 1/6 + 1/3 * 2 = 5/6$
- f_2: price = 2, $CS = 2/3$
- f_3: price = 2, $CS = 0$
- Consumer surplus $5/6 * 2/3 + 2/3 * 1/6 = 2/3$
- Other way to see: posting the lowest price in the support of each posterior distribution (hence efficient allocation) is optimal for each distribution

How to visualize this? Consider the set of all distributions.

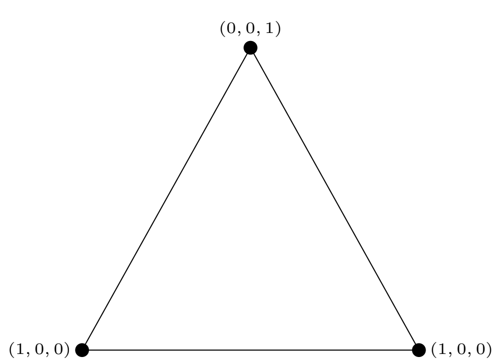

Divide this set into three sets:

- Price 1 best: $1 \geq 2(f(2)+f(3)), 3f(3)$
- Price 2 best $2(f(2)+f(3)) \geq 1, 3f(3)$
- Price 3 best: $3f(3) \geq 1, 2(f(2)+f(3))$

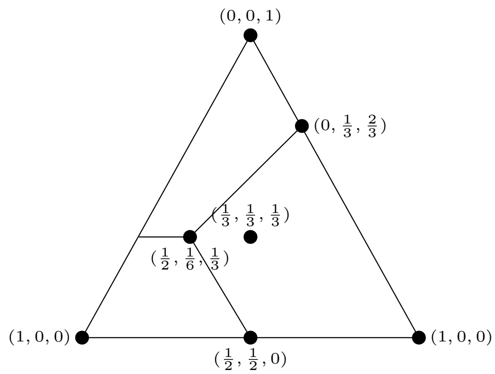

Since the prior distribution $(\frac{1}{3},\frac{1}{3},\frac{1}{3})$ is inside the shape defined by vertices $(\frac{1}{2},\frac{1}{6},\frac{1}{3})$, $(0,\frac{1}{3},\frac{2}{3})$, $(0,1,0)$, and $(\frac{1}{2},\frac{1}{2},0)$, it can be written as a weighted average of those vertices. In fact, since the prior is inside the triangle defined by vertices $(\frac{1}{2},\frac{1}{6},\frac{1}{3})$, $(0,\frac{1}{3},\frac{2}{3})$, and $(\frac{1}{2},\frac{1}{2},0)$, the prior can be written as a weighted average of those vertices, as we did with the test that we introduced before. The optimal test (for the consumer) is not unique as the prior can also be written as a weighted average of posteriors $(\frac{1}{2},\frac{1}{6},\frac{1}{3})$, $(0,\frac{1}{3},\frac{2}{3})$, and $(0,1,0)$.

3 Buyer optimal learning

This is form [Roesler and Szentes, 2017].

In [Bergemann et al., 2015], the seller receives information about the buyer's value. Here we think about the case where the buyer receives information about her own value.

Standard setup:

- Single seller, single buyer, single product
- Buyer's value is $\theta \in [0,1]$. Neither the seller nor buyer know θ.
- As in our dynamic mechanism design setup, the buyer observes a signal s that is informative of θ. I.e., there is a joint distribution μ over (s,θ). Both seller and buyer know μ, and the buyer privately observes s.
- Let F_θ and F_s denote marginals on θ and s.

Knowing μ, the seller chooses a price p to maximize revenue. Given a price p, the buyer's expected surplus from buying is $E[\theta|s] - p$ when her signal is s. So she buys if and only if $E[\theta|s] \geq p$. So the seller's problem is

$$\max_p pPr_{s\sim F_s}[E[\theta|s] \geq p].$$

Given a price p, the buyer's expected surplus is

$$E_{s\sim F_s}[(E[\theta|s] - p)\mathbb{1}_{E[\theta|s]\geq p}].$$

Here is the question: suppose that the buyer's distribution over values is fixed, but she could choose how much and/or what to "learn" about her value. Formally, suppose that she could choose any μ such that $F_\theta = F$ for some given F. Then what μ would she choose to maximize her surplus?

Since F_θ is fixed, we can think of choosing μ as choosing a mapping $\pi : \Theta \to \Delta(S)$ (which we called an "investigation" in [Kamenica and Gentzkow, 2011]).

Example: Suppose that F is uniform over $\{1, 2\}$. Suppose that π is as follows.

$\pi(s\|\theta)$	L	H
1	q	$1-q$
2	$1-q$	q

For some $q \in [0.5, 1]$. Notice that choosing π is equivalent to choosing a joint distribution μ

(θ, s)	(1,L)	(1,H)	(2,L)	(2,H)
$Pr[(\theta, s)]$	$q/2$	$(1-q)/2$	$(1-q)/2$	$q/2$

with marginal over values F.

What is the optimal price?

$$\max_p pPr_{s\sim F_s}[E[\theta|s] \geq p].$$

To answer this question, we have to think about $E[\theta|s]$. We have $E[\theta|L] = 2 - q$ and $E[\theta|H] = 1 + q$. So $E[\theta|s]$ is equal to $2 - q$ with probability $1/2$ (when $s = L$), and is equal to $1 + q$ with probability $1/2$ (when $s = H$).

So revenue is $2-q$ if price is $2-q$, and is $(1+q)/2$ if $p = 1+q$. It is optimal to set $p = 2-q$, in which case the buyer's surplus is $(1 + q)/4$. So to maximize surplus, the buyer chooses $q = 1$, i.e., to learn her value perfectly.

Boring.

But this should be obvious in hindsight. The seller never sets a price less than 1. So the surplus of the buyer is at most the expected value, 1.5, minus 1. This is accomplished with full information.

Now repeat the exercise when the prior is uniform over $\{1, 3\}$. Optimize over the same set of investigations π. Convince yourself that with perfect learning, $q = 1$, or with no learning at all, $q = 1/2$, the consumer surplus is zero. But the best learning is when $q = 5/6$, in which case buyer surplus is positive (1/6). Interesting! Best leaning happens to "maximize learning" subject to the condition that the buyer buys with probability 1. This happens to be the case in general, i.e., when maximizing over all investigations.

Notice that once π (or μ) is fixed, the problem becomes the standard screening problem, where the buyer's value is either $2 - q$ or $1 + q$ each with probability 1/2, the buyer "knows her value", but the seller knows only the distribution. The value in this transformed problem is just the expected θ for a given s. To summarize, we can see that for any π (or μ), there exists a distribution F^π (corresponding to the distribution of expected values) such that the seller chooses a p to maximize the expected revenue

$$\max_p pPr_{v\sim F^\pi}[v \geq p].$$

One annoying thing is that F^π may have mass points. So the probability of purchase is not $1 - F^\pi(p)$, which would be the probability that $v < p$. So let's invent notation $D(F^\pi, p) = Pr_{v\sim F^\pi}[v \geq p]$ and then we can write revenue as

$$pD(F^\pi, p)$$

and the buyer's surplus is

$$E_{v\sim F^\pi}[(v - p)\mathbb{1}_{v\geq p}].$$

Let $\mathcal{F} = \{F^\pi|\pi\}$ be the set of all possible distributions over expected values. The buyer's problem can be written as

$$\max_{p,G\in\mathcal{F}} E_{v\sim G}[(v - p)\mathbb{1}_{v\geq p}]$$
$$\text{s.t. } p \in \arg\max pD(G, p).$$

Lemma 3.1 *$G \in \mathcal{F}$ if and only if G is a "mean-preserving contraction of F, that is*

$$\int_0^x G(v)dv \leq \int_0^x F(v)dv,$$

with equality at $x = 1$.

Proof: This is standard (see e.g., [Mas-Colell et al., 1995] Definition 6.D.2). **QED**

As a sanity check, verify that the uniform distribution over $2 - q$ and $1 + q$ satisfies this condition.

By integration by parts, the equality $\int_0^1 G(v)dv = \int_0^1 F(v)dv$ is the same as $\int_0^1 vdG(v) = \int_0^1 vdF(v)$. That is, the distributions must have the same mean.

Now we define a class of distributions and show that the optimal distribution belongs to that class. In particular, for $q, B \in [0, 1], q \leq B$, define

$$G_q^B(v) = \begin{cases} 0 & \text{if } v \in [0, q) \\ 1 - q/v & \text{if } v \in [q, B] \\ 1 & \text{if } v \in [B, 1] \end{cases}$$

The important property is that the seller is indifferent between choosing any price in $[q, B]$.

Proposition 3.2 *There exists (q, B) such that G_q^B and price $p = q$ are optimal.*

Proof: Consider any feasible solution G with revenue-maximizing price p. By Lemma 3.1, G must be a MPC of F, i.e., $\int_0^x G(v)dv \leq \int_0^x F(v)dv$ for all x, with equality at $x = 1$. Let π be seller's profit.

First, we show that there exists B such that G_π^B is a MPC of F. We do so by showing that G_π^B is a MPC of G.

B is constructed so that G_π^B has the same expectation as G. Consider two cases.

(1) $B = \pi$. then the expectation of G_π^B is π. Since π is the optimal profit for G, it must be at most the expected value under G. So

$$\int_0^1 v dG_\pi^\pi(v) \leq \int_0^1 v dG(v),$$

which means that

$$\int_0^1 G_\pi^\pi(v)dv \geq \int_0^1 G(v)dv, \tag{1}$$

(2) $B = 1$. Since π is the optimal profit for G, we have $v(1 - G(v)) \leq \pi$ for each $v \geq \pi$, so $G(v) \geq 1 - \pi/v = G_\pi^1(v)$, and so the area under G is weakly higher than the area under G_π^1,

$$\int_0^1 G_\pi^1(v)dv \leq \int_0^1 G(v)dv \tag{2}$$

Given (1) and (2) and by the intermediate value theorem, there exists B such that the area under G is equal to the area under G_π^B.

We argued above that for any $v \leq B$, we have $G_\pi^B(v) \leq G(v)$. So $\int_0^x G_\pi^B(v)dv \leq \int_0^x G(v)dv$. Also for $v \geq B$, we have $G_\pi^B(v) = 1 \geq G(v)$. So

$$\int_0^x G_\pi^B(v)dv = \int_0^1 G_\pi^B(v)dv - \int_x^1 G_\pi^B(v)dv \leq \int_0^1 G(v)dv - \int_x^1 G(v)dv = \int_0^x G(v)dv.$$

So we have established that G_π^B is an MPC of F.

Price π maximizes revenue for G_π^B. So it only remains to show that expected buyer surplus for (π, G_π^B) is weakly higher than that for (p, G). The consumer surplus with (π, G_π^B) is $\mu - \pi$. The consumer surplus with (G, p) is at most $\mu - \pi$ since the sum of producer and consumer surplus is at most the expected value μ. **QED**

Notice that when $p = q$ and the distribution of values is G_q^B, then the buyer buys with probability 1. So her payoff is $E_{v \in G_q^B} - p$, which by the MPC constraint is $E_{\theta \sim F}[\theta] - p$. So the highest surplus for the buyer is obtained by choosing $q^* = \min_q \text{s.t. } \exists B, G_q^B \in \mathcal{F}$.

Example: θ is uniform on $[0, 1]$. Then G_q^B with $q \simeq 0.2$ and $B \simeq 0.87$ is optimal. Producer surplus is q, and consumer surplus is $E[\theta] - q \simeq 0.3$. Compare this with full information, where producer surplus is 0.25, and consumer surplus is 0.125. Deadweight loss is 0.125. Notice that producer surplus is lower with optimal learning than with full information.

The above observation is general: For any $G \in \mathcal{F}$, $pD(p,G) \geq q^*$. Define $\pi = pD(p,G)$. The lemma above showed that G^B_π is in $\mathcal{F}$ for some B. Since $q^* = \min_q$ s.t. $\exists B, G^B_q \in \mathcal{F}$, we have $q^* \leq \pi$.

So we can now characterize the set of consumer surplus and producer surplus pairs for all information structures. It is a triangle given by $(0,\mu),(0,q^*),(\mu - q^*, q^*)$.

References

[Bergemann et al., 2015] Bergemann, D., Brooks, B., and Morris, S. (2015). The limits of price discrimination. American Economic Review, 105(3):921–57.

[Kamenica and Gentzkow, 2011] Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. American Economic Review, 101(6):2590–2615.

[Mas-Colell et al., 1995] Mas-Colell, A., Whinston, M. D., Green, J. R., et al. (1995). Microeconomic theory, volume 1. Oxford university press New York.

[Roesler and Szentes, 2017] Roesler, A.-K. and Szentes, B. (2017). Buyer-optimal learning and monopoly pricing. American Economic Review, 107(7):2072–80.