

IM 7011: Information Economics (Fall 2014)

The Signaling Theory

Ling-Chieh Kung

Department of Information Management
National Taiwan University

November 24, 2014

Road map

- ▶ **Introduction.**
- ▶ Bayesian updating.
- ▶ The first example.

Signaling

- ▶ We have studied two kinds of principal-agent relationship:
 - ▶ Screening: the agent has hidden information.
 - ▶ Moral hazard: the agent has hidden actions.
- ▶ Starting from now, we will study the third situation: **signaling**.
 - ▶ The **principal** will have hidden information.
- ▶ Both screening and signaling are adverse selection issues.

Origin of the signaling theory

- ▶ Akerlof (1970) studies the market of **used cars**.
 - ▶ The owner of a used car knows the **quality** of the car.
 - ▶ Potential buyers, however, do not know it.
 - ▶ The quality is hidden information observed only by the principal (seller).
- ▶ What is the issue?
 - ▶ Buyers do not want to buy “lemons”.
 - ▶ They only pay a price for a used car that is “**around average**”.
 - ▶ Owners of **bad** used cars are happy for selling their used cars.
 - ▶ Owners of **good** ones do not sell theirs.
 - ▶ Days after days... there are only bad cars on the market.
 - ▶ The “expected quality” and “average quality” become lower and lower.
- ▶ **Information asymmetry** causes **inefficiency**.
 - ▶ In screening problems, information asymmetry protects agents.
 - ▶ In signaling problems, information asymmetry **hurts everyone**.
- ▶ That is why we need platforms that suggest prices for used cars.

Origin of the signaling theory

- ▶ Spence (1973) studies the market of **labors**.
 - ▶ One knows her **ability** (productivity) while potential employers do not.
 - ▶ The “quality” of the worker is hidden.
 - ▶ Firms only pay a wage for “**around average**” workers.
 - ▶ **Low**-productivity workers are happy. **High**-productivity ones are sad.
 - ▶ Productive workers leave the market (e.g., go abroad). Wages decrease.
- ▶ What should we do? No platform can suggest wages for individuals!
- ▶ That is why we get **high education** (or study in good schools).
 - ▶ It is not very costly for a high-productivity person to get a higher degree.
 - ▶ It is **more costly** for a low-productivity one to get it.
 - ▶ By getting a higher degree (e.g., a master), high-productivity people **differentiate** themselves from low-productivity ones.
 - ▶ Getting a higher degree is **sending a signal**.
- ▶ This will happen (as an equilibrium) even if education itself **does not** enhance productivity!
 - ▶ Though this may not be a good thing, it seems to be true.
 - ▶ Think about **certificates**.

Signaling

- ▶ Signaling is for the principal to send a message to the agent to **signal the hidden information**.
 - ▶ Sending a message requires an **action** (e.g., getting a degree).
- ▶ For signaling to be effective, different types of principal should take different actions.
 - ▶ It must be **too costly** for a type to take a certain action.
- ▶ Other examples:
 - ▶ A manufacturer offers a **warranty** policy to signal the product reliability.
 - ▶ A firm sets a high **price** to signal the product quality.
 - ▶ “Full **refund** if not tasty”.

Signaling games

- ▶ How to model and analyze a signaling game?
 - ▶ There is a principal and an agent.
 - ▶ The principal has a **hidden type**.
 - ▶ The agent cannot observe the type and thus have a **prior belief** on the principal's type.
 - ▶ The principal chooses an **action** that is observable.
 - ▶ The agent then forms a **posterior belief** on the type.
 - ▶ Based on the posterior belief, the agent **responds** to the principal.
- ▶ The principal takes the action to **alter** the agent's belief.
- ▶ An example:
 - ▶ A firm makes and sells a product to consumers.
 - ▶ The **reliability** of the product is hidden.
 - ▶ Consumers have a prior belief on the reliability.
 - ▶ The firm chooses between **offering a warranty or not**.
 - ▶ By observing the policy, the consumer **updates his belief** and make the purchasing decision accordingly.
- ▶ We need to model belief updating by **the Bayes' theorem**.

Road map

- ▶ Introduction.
- ▶ **Bayesian updating.**
- ▶ The first example.

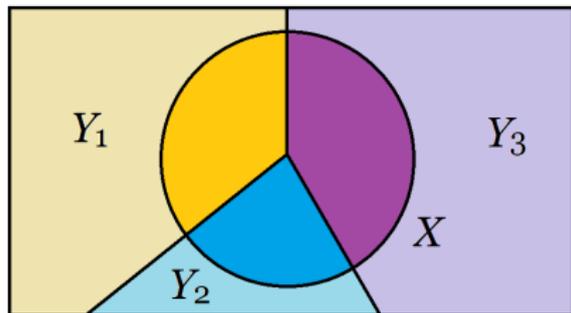
Law of total probability

- ▶ The following law is a component of Bayes' rule:

Proposition 1 (Law of total probability)

Let events Y_1, Y_2, \dots , and Y_k be mutually exclusive and completely exhaustive and X be another event, then

$$\Pr(X) = \sum_{i=1}^k \Pr(Y_i) \Pr(X|Y_i).$$



Belief updating

- ▶ For some unknowns, we have some original estimates.
- ▶ We form a **prior belief** or assign a **prior probability** to the occurrence of an event.
 - ▶ Before I toss a coin, my belief of getting a head is $\frac{1}{2}$.
- ▶ If our estimation is accurate, the **relative frequency** of the occurrence of the event should be **close** to my prior belief.
 - ▶ In 100 trials, probably I will see 48 heads. $\frac{48}{100} \approx \frac{1}{2}$.
 - ▶ What if I see 60 heads? What if 90?
- ▶ In general, we expect observations to follow our prior belief.
- ▶ If this is not the case, we probably should update our prior belief into a **posterior belief**.

Example: Popularity of a product

- ▶ Suppose we have a product to sell.
- ▶ We do not know how consumers like it.
- ▶ Two possibilities (events): popular (P) and unpopular (U).
 - ▶ Our **prior** belief on P is 0.7.
 - ▶ We believe, with a 70% probability, that the product is popular.
- ▶ When one consumer comes, she may buy it (B) or go away (G).
 - ▶ If popular, the buying probability is 0.6.
 - ▶ If unpopular, the buying probability is 0.2.
- ▶ Suppose event G occurs once, what is our **posterior** belief?

Example: Popularity of a product

- ▶ We have the marginal probabilities $\Pr(P)$ and $\Pr(U)$:

	<i>B</i>	<i>G</i>	Total
<i>P</i>	?	?	0.7
<i>U</i>	?	?	0.3
Total	?	?	1

- ▶ We have the conditional probabilities:
 - ▶ $\Pr(B|P) = 0.6 = 1 - \Pr(G|P)$ and $\Pr(B|U) = 0.2 = 1 - \Pr(G|U)$.
- ▶ We thus can calculate those joint probabilities:

	<i>B</i>	<i>G</i>	Total
<i>P</i>	0.42	0.28	0.7
<i>U</i>	0.06	0.24	0.3
Total	?	?	1

Example: Popularity of a product

- ▶ We now can calculate the marginal probabilities $\Pr(B)$ and $\Pr(G)$:

	<i>B</i>	<i>G</i>	Total
<i>P</i>	0.42	0.28	0.7
<i>U</i>	0.06	0.24	0.3
Total	0.48	0.52	1

- ▶ Now, we observe one consumer going away (event G).
- ▶ What is the posterior belief that the product is popular (event P)?
 - ▶ This is the conditional probability $\Pr(P|G) = \frac{\Pr(P \cap G)}{\Pr(G)} = \frac{0.28}{0.52} \approx 0.54$.
- ▶ Note that we **update our belief** on P from 0.7 to 0.54.
- ▶ The fact that one goes away makes us **less confident**.
- ▶ If another consumer goes away, the updated belief on P becomes 0.37.
 - ▶ Use the old posterior as the new prior.
 - ▶ Use $\Pr(P|G)$ as $\Pr(P)$ and $\Pr(U|G)$ as $\Pr(U)$ and repeat.
- ▶ After five consumers go away in a row, the posterior becomes 0.07.
 - ▶ We tend to believe the product is unpopular!

Bayes' theorem

- ▶ By the law of total probability, we establish **Bayes' theorem**:

Proposition 2 (Bayes' theorem)

Let events Y_1, Y_2, \dots , and Y_k be mutually exclusive and completely exhaustive and X be another event, then

$$\Pr(Y_j|X) = \frac{\Pr(Y_j \cap X)}{\Pr(X)} = \frac{\Pr(Y_j) \Pr(X|Y_j)}{\sum_{i=1}^k \Pr(Y_i) \Pr(X|Y_i)} \quad \forall j = 1, 2, \dots, k.$$

- ▶ Sometimes we have events $\{Y_i\}_{i=1, \dots, k}$ and X :
 - ▶ It is clear how Y_i s affect X but not the other way.
 - ▶ Bayes' theorem is applied to **use X to infer $\{Y_i\}_{i=1, \dots, k}$** .
- ▶ P and U naturally affect G and B but not the other way.
 - ▶ So we apply Bayes' theorem to use G to infer P and U :

$$\Pr(P|G) = \frac{\Pr(P) \Pr(G|P)}{\Pr(P) \Pr(G|P) + \Pr(U) \Pr(G|U)} = \frac{0.7 \times 0.4}{0.7 \times 0.4 + 0.3 \times 0.8} = 0.54.$$

Road map

- ▶ Introduction.
- ▶ Bayesian updating.
- ▶ **The first example.**

The first example

- ▶ A firm makes and sells a product with **hidden reliability** $r \in (0, 1)$.
 - ▶ r is the probability for the product to be functional.
- ▶ If a consumer buys the product at price t :
 - ▶ If the product works, his utility is $\theta - t$.
 - ▶ If the product fails, his utility is $-t$.
- ▶ The firm may offer a **warranty** plan and repair a broken product.
 - ▶ The firm pays the repairing cost $k > 0$.
 - ▶ The consumer's utility is $\eta \in (0, \theta)$.
- ▶ The price is fixed (exogenous).
- ▶ Suppose $w = 1$ if a warranty is offered and 0 otherwise.
- ▶ Expected utilities:
 - ▶ The firm's expected utility is $u_F = t - (1 - r)kw$.
 - ▶ The consumer's expected utility is $u_C = r\theta + (1 - r)\eta w - t$.
- ▶ The consumer buys the product if and only if $u_C \geq 0$.
- ▶ The firm chooses whether to offer the warranty accordingly.

The first example: no signaling

- ▶ Suppose $r \in \{r_H, r_L\}$: The product may be reliable or unreliable.
 - ▶ $0 < r_L < r_H < 1$.
- ▶ Under complete information, the decisions are simple.
 - ▶ The firm's expected utility is $u_F = t - (1 - r_i)kw$.
 - ▶ The consumer's expected utility is $u_C = r_i\theta + (1 - r_i)\eta w - t$.
- ▶ Under incomplete information, they may make decision according to the **expected reliability**:
 - ▶ Let $\beta = \Pr(r = r_L) = 1 - \Pr(r = r_H)$ be the consumer's **prior belief**.
 - ▶ The expected reliability is $\bar{r} = \beta r_L + (1 - \beta)r_H$.
 - ▶ The firm's expected utility is $u_F = t - (1 - r_i)kw$.
 - ▶ The consumer's expected utility is $u_C = \bar{r}\theta + (1 - \bar{r})\eta w - t$.
- ▶ But wait! The **unreliable** firm will tend to offer **no warranty**.
 - ▶ Because $(1 - r_L)k$ is high.
 - ▶ This forms the basis of **signaling**.

The first example: signaling

- ▶ Below we will work with the following parameters:

- ▶ $r_L = 0.2$ and $r_H = 0.8$.
- ▶ $\theta = 20$ and $\eta = 5$.
- ▶ $t = 11$ and $k = 15$.

- ▶ Payoff matrices (though players make decisions sequentially):

		Consumer	
		Buy	Not
Firm	$w = 1$	8, 6	0, 0
	$w = 0$	11, 5	0, 0

(Product is reliable)

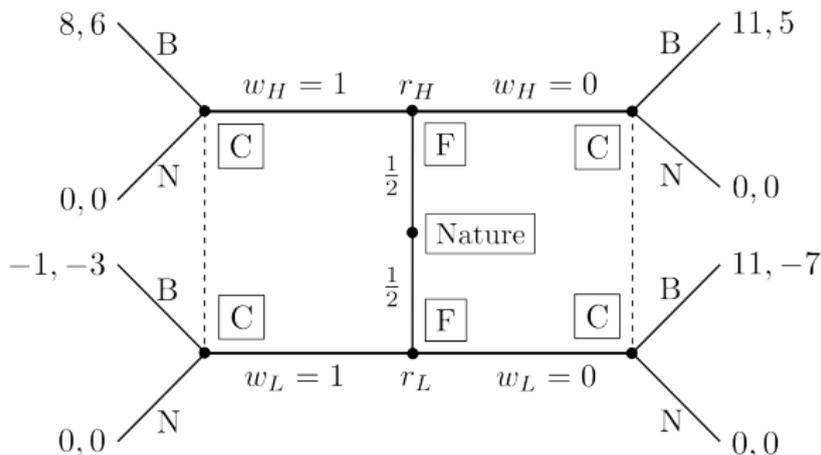
		Consumer	
		Buy	Not
Firm	$w = 1$	-1, -3	0, 0
	$w = 0$	11, -7	0, 0

(Product is unreliable)

- ▶ The issue is: The consumer does not know **which matrix** he is facing!
- ▶ The reliable firm tries to convince the consumer that it is the first one.

Game tree

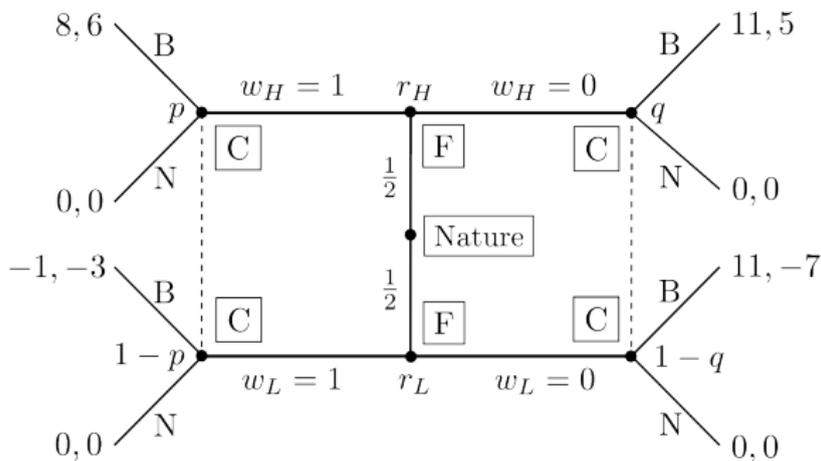
- ▶ We express this **game with incomplete information** by the following game tree:
 - ▶ $\boxed{\text{F}}$ and $\boxed{\text{C}}$: players.
 - ▶ $\boxed{\text{Nature}}$: a fictitious player that draws the type randomly.
 - ▶ Let $\beta = \frac{1}{2}$ be the prior belief.



Concept of equilibrium

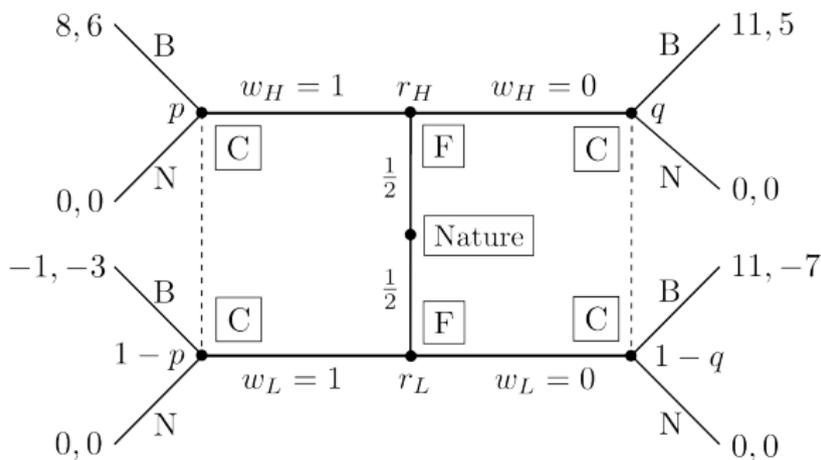
- ▶ What is a (pure-strategy) **equilibrium** in a signaling game?
 - ▶ Decisions:
 - ▶ The “two” firms’ actions: (w_H, w_L) , $w_i \in \{0, 1\}$.
 - ▶ The consumer’s strategy: (a_1, a_0) , $a_j \in \{B, N\}$.
 - ▶ Posterior beliefs:
 - ▶ Let $p = \Pr(r_H | w = 1)$ be the posterior belief upon observing a warranty.
 - ▶ Let $q = \Pr(r_H | w = 0)$ be the posterior belief upon observing no warranty.
 - ▶ An equilibrium is a strategy-belief **profile** $((w_H, w_L), (a_1, a_0), (p, q))$:
 - ▶ No firm wants to deviate based on the consumer’s posterior belief.
 - ▶ The consumer does not deviate based on his posterior belief.
 - ▶ The beliefs are updated according to the firms’ actions by the Bayes’ rule.
 - ▶ It is extremely hard to “search for” an equilibrium. It is easier to “**check**” whether a given profile is one.
 - ▶ We start from the firms’ actions:¹
 - ▶ Can $(1, 0)$ be part of an equilibrium? How about $(0, 1)$, $(1, 1)$, and $(0, 0)$?
- ¹It is typical to start from the principal’s actions.

Warranty for the reliable product only



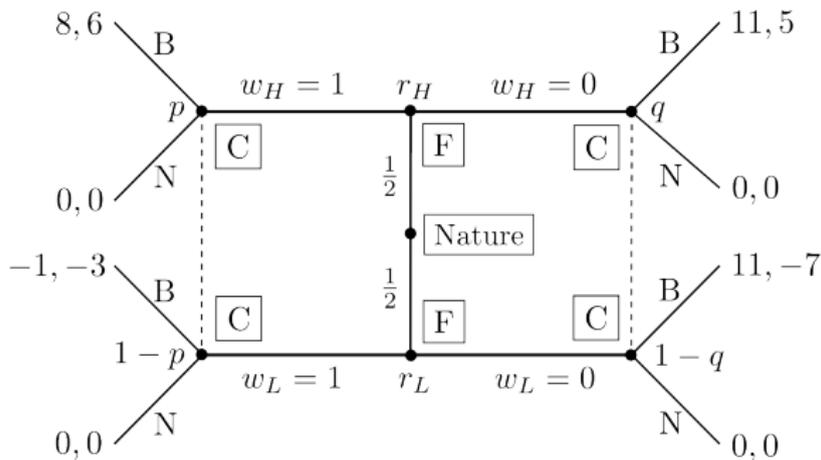
- ▶ We start from $((1, 0), (a_1, a_0), (p, q))$.
- ▶ Bayesian updating: $p = 1, q = 0$: $((1, 0), (a_1, a_0), (1, 0))$.
- ▶ Consumer $((1, 0), (B, N), (1, 0))$.
- ▶ No firm wants to deviate.

Warranty for the unreliable product only



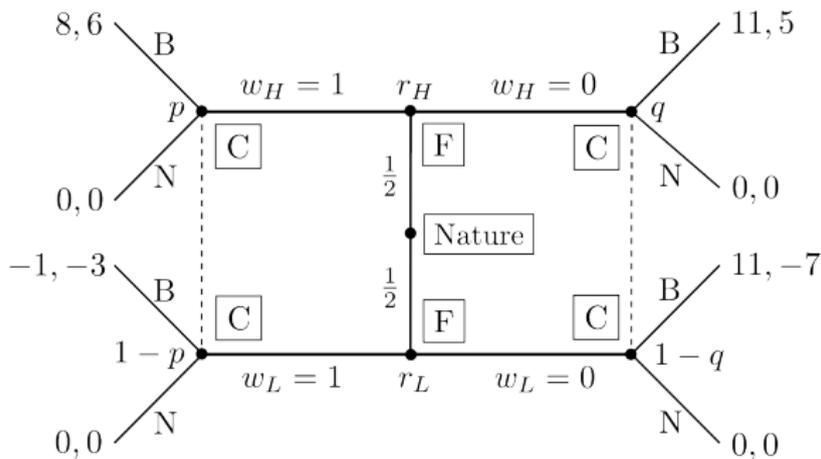
- ▶ We start from $((0, 1), (a_1, a_0), (p, q))$.
- ▶ Bayesian updating: $p = 0, q = 1$: $((0, 1), (a_1, a_0), (0, 1))$.
- ▶ Consumer: $((0, 1), (N, B), (0, 1))$.
- ▶ But now the unreliable firm deviates to $w_L = 0$!

Both offering warranties



- ▶ We start from $((1, 1), (a_1, a_0), (p, q))$.
- ▶ Bayesian updating: $p = \frac{1}{2}$, $q \in [0, 1]$: $((1, 1), (a_1, a_0), (\frac{1}{2}, [0, 1]))$.
- ▶ Consumer: $((1, 1), (B, \{B, N\}), (\frac{1}{2}, [0, 1]))$.
- ▶ If $a_0 = B$, no firm offers a warranty: $((1, 1), (B, N), (\frac{1}{2}, [0, 1]))$.
- ▶ But now the unreliable firm deviates to $w_L = 0!$

Both offering no warranty



- ▶ We start from $((0, 0), (a_1, a_0), (p, q))$.
- ▶ Bayesian updating: $p \in [0, 1]$, $q = \frac{1}{2}$: $((0, 0), (a_1, a_0), ([0, 1], \frac{1}{2}))$.
- ▶ Consumer: $((0, 0), (B, N), ([\frac{1}{3}, 1], \frac{1}{2}))$, or $((0, 0), (N, N), ([0, \frac{1}{3}], \frac{1}{2}))$.
- ▶ For the former, the reliable firm deviates to $w_H = 1$. The latter is a pooling equilibrium.

Interpretations

- ▶ There are **pooling**, **separating**, and **semi-separating** equilibria:
 - ▶ In a pooling equilibrium, all types take the same action.
 - ▶ In a separating equilibrium, different types take different actions.
 - ▶ In a semi-separating one, some but not all types take the same action.
- ▶ In this example, there are two (sets of) equilibria:
 - ▶ A separating equilibrium $((1, 0), (B, N), (1, 0))$.
 - ▶ A pooling equilibrium $((0, 0), (N, N), ([0, \frac{1}{3}], \frac{1}{2}))$.
- ▶ What does that mean?

Interpretations

- ▶ The separating equilibrium is $((1, 0), (B, N), (1, 0))$:
 - ▶ The reliable product is sold with a warranty.
 - ▶ The unreliable product, offered with no warranty, is not sold.
 - ▶ The reliable firm **successfully signals** her reliability.
 - ▶ The system becomes more efficient.
 - ▶ Because it is too costly for the unreliable firm to do the same thing.
- ▶ The pooling equilibrium is $((0, 0), (N, N), ([0, \frac{1}{3}], \frac{1}{2}))$.
 - ▶ Both firms do not offer a warranty.
 - ▶ The consumer cannot update his belief.
 - ▶ The consumer does not buy the product.
- ▶ In this (and most) signaling game, there are **multiple** equilibria.