Contents

Chapter 11: Signaling and Screening

How can we overcome the market failure due to asymmetric information? This chapter considers the two important methods known as signaling and screening.

In a signaling model, the informed party sends a signal to the uninformed party. However, it is often the case that there is an advantage to pretending to be something you are not. In equilibrium in a successful signaling model, each type sends a different signal so the uniformed party can infer type via the signal. This is called a separating equilibrium. If all types send the same signal then the signal does not convey any credible information. In this case we have an unsuccessful signaling model with a pooling equilibrium. Either way, we check the PBE property that beliefs are consistent with Bayes' rule and the equilibrium strategies. We show this for a variant of Spence's model, where the worker's education level is a signal to a potential employer who is uniformed about an individual worker's ability or type.

We then turn to screening models where the uniformed party sets up a menu of options and the informed party's choice can credibly reveal its type. We illustrate with a model of insurance where high risk drivers choose a policy with a high price and a low deductible.

1 Methods of Dealing With Asymmetric Information

How can the asymmetric information market failure be partially or completely overcome?

1. Destroy the information. If sellers don't know the type (θ) then all types are supplied since sellers participation constraint can't compare the actual type and hence reservation value with the price. Example: bags of wholesale diamonds (or oranges) vs selections handpicked by the vendor.

2. Destroy the choice. Don't allow the informed party to leave the market. Again, this eliminates the participation constraint. For example, if a government insurance

program is mandatory, then there is no "shrinking pool" with the resulting adverse selection. This is known as the mandate in the ACA (aka Obamacare).

3. Publicize the information. This can be done by:

   (a) Third party inspection. For example, pay a mechanic to check the car out and determine whether it is a lemon. Pay Moody's or Standard & Poor's to rate the quality of an asset like a bond, or corporate debt.

   (b) Signaling. Let the informed party convey credible information by sending a signal to the uninformed party.

   (c) Screening. The uniformed party offers a menu of choices to the informed party. The chosen option credibly reveals their type.

      - Ex: Insurance company offers two policies:

      i. Low premium, high deductible or

      ii. High premium, low deductible.

      If the menu of choices is properly designed then the high risk types will choose (ii) and the low risk types will choose (i).

2 Spence's Signaling Model

Michael Spence (1973, 1974) developed a pure "sheepskin effect" model where education only sends a signal. The second paper shows the main results hold also when education increases productivity.

- Job market signaling.

  – The informed party is the worker who chooses an education level.

  – Education is a message or a signal about type sent by the informed party to the uniformed party.

  - The uniformed party observes the message and then updates their beliefs about the sender's type and then chooses an action.
  - Payoffs are determined by type, message and action.

- Set of types denoted $\Theta = \{\theta_H, \theta_L\}$ where θ_H is high productivity and θ_L is low productivity workers.
  - The prior distribution of types is common knowledge.

- Nature chooses type $\theta \in \Theta$ of sender (informed player).

- The sender (informed type) chooses a message $m(\theta) \in M$ from the set of possible messages. To get interesting results, the cost of sending at least one particular signal must differ across types. Here messages are years of education, and the cost is lower for more productive types.

- The respondent (uniformed player) observes the message and then chooses an action $a(m) \in A$ from the set of allowable actions, after forming beliefs $\mu(\theta|m)$ about the sender's type. The beliefs are a Bayesian posterior probability. The action may be to hire or not hire a potential worker.

3 Types of Equilibria

The equilibrium concept is a perfect Bayesian Nash equilibrium (PBE) of several possible types: pooling, separating, partial pooling or hybrid.

1. Pooling.
   - All sender types S_θ choose the same $m \in M$.
   - This is an unsuccessful signal.

- The respondent does not update her priors regarding types.

2. Separating.

   - Each S_θ chooses a different m.
   - This is a successful signal.
     - The signal perfectly reveals the type.
     - The Bayesian posterior is {0,1}.
     - This requires at least as many signals as types.

3. Partial pooling.

   - Somewhere in between.
   - More than one type chooses the same message, but more than one message is chosen across types.
   - Here the Bayesian posterior is more accurate (better) than the prior, but not fully revealing.

4. Hybrid.

   - Some type chooses a given m with probability 1 and at least one other type chooses that same message with a positive probability not equal to one.
   - Again, the Bayesian posterior is better than the prior, but not fully revealing.

- Example: suppose three types $\Theta = \{\theta_H, \theta_M, \theta_L\}$ and two messages $M = \{yes, no\}$.

1. If the θ_H and θ_M types choose *yes* and the θ_L types choose *no*, then there is partial pooling.

   (a) - If *yes* is the message, the respondent knows the Bayesian posterior $\mu(\theta_H|yes)$ is greater than the prior $p(\theta_H)$, but less than 1.

- The Bayesian posterior $\mu(\theta_M|yes)$ is greater than the prior $p(\theta_M)$, but less than 1.
- The Bayesian posterior $\mu(\theta_L|no) = 1$ and $\mu(\theta_L|yes) = 0$.

2. If all the θ_H types choose *yes* and only some of the θ_M types choose *yes*, then we have a hybrid equilibrium.

   (a)
   - Again, the Bayesian posterior $\mu(\theta_H|yes)$ is greater than the prior $p(\theta_H)$, but less than 1.
   - And again the Bayesian posterior $\mu(\theta_L|no) = 1$.
   - Here, however, $\mu(\theta_M|yes)$ could be greater or less than the prior $p(\theta_M)$.

4 Signaling Game Example

- The simplest non-trivial case is: two types, two messages and two actions.

$$\Theta = \{\theta_H, \theta_L\},\ M = \{r, l\},\ A = \{u, d\} \tag{1}$$

- Nature chooses $\theta \in \Theta$, the sender (informed party) chooses $m(\theta) \in M$ and the respondent (uniformed party) chooses $a(m) \in A$.
- The strategy sets are complete contingency plans which are mappings from the information sets to the action sets.
- The sender has type (θ) in their information set and chooses message (m).
- The respondent has message (m) in their information set and chooses action (a).

- Sender has four pure strategies:

$$\begin{aligned} s_1 &= (r|\theta_H, r|\theta_L) \\ s_2 &= (r|\theta_H, l|\theta_L) \\ s_3 &= (l|\theta_H, r|\theta_L) \\ s_4 &= (l|\theta_H, l|\theta_L) \end{aligned} \tag{2}$$

- s_1 and s_4 are pooling strategies since both types choose the same message.
  - In this case the signal is unsuccessful in overcoming the asymmetric information.
- Strategies s_2 and s_3 are separating since each type chooses a different message.
  - Here the signal fully reveals the type and the signal is successful in overcoming the asymmetric information.
- Thus in this simple signalling game, the senders' strategies determine the type of equilibrium.
- Responder also has four pure strategies:

$$\begin{aligned} r_1 &= (u|r, u|l) \\ r_2 &= (u|r, d|l) \\ r_3 &= (d|r, u|l) \\ r_4 &= (d|r, d|l) \end{aligned} \tag{3}$$

- We could write a 4x4 NFG (bimatrix) and use a brute force approach.
  - This can be laborious.

– Often it is better to use the "guess and check" method. You can guess that there is a separating PBE, and check a few conditions to confirm or to eliminate this possibility. Likewise for a pooling PBE.

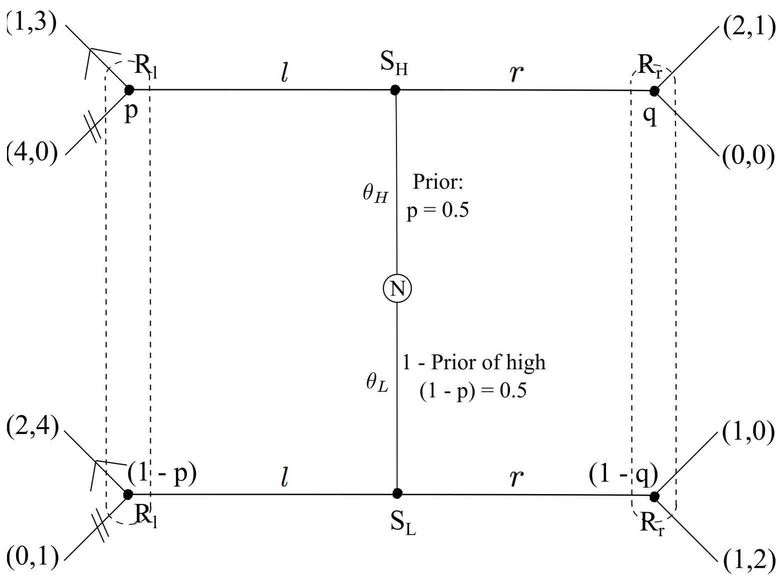

Figure 1: Example from Gibbons, pg. 189. The payoffs are for sender, then responder.

- After guessing the Sender's strategy, we compute the Responder's BRs.
  – Let p be Responder's posterior probability for θ_H when Sender chooses l.
  – If the sender chooses l in Figure 1, then:

$$a^*(l) = u \quad \forall p \in [0,1] \tag{4}$$

  – Let q be Responder's posterior probability for θ_H when Sender chooses r.
  – If the Sender chooses r then Responder's expected payoffs are:

$$\begin{aligned} E\pi_R(u|r) &= q \\ E\pi_R(d|r) &= 2(1-q) \\ E\pi_R(u|r) - E\pi_R(d|r) &= 3q-2 \\ 3q-2 &> 0 \quad \forall q > \frac{2}{3} \end{aligned} \tag{5}$$

– Thus:

$$a^*(r) = \left\{ \begin{array}{ll} d & \forall\ q \leq \frac{2}{3} \\ u & \forall\ q \geq \frac{2}{3} \end{array} \right\} \tag{6}$$

5 Solving Signalling Games

Some definitions will be helpful.

- *Equilibrium path*: In an EFG an information set is on the equilibrium path if there is a positive probability that this information set will reached if the game is played according to the equilibrium strategies.
  - An information set is off the equilibrium path if there is zero probability that it will be reached.
- A perfect Bayesian-Nash equilibrium (PBE) requires that beliefs are determined by Bayes' rule and the players equilibrium strategies at information sets along the equilibrium path.
  - At information sets off the equilibrium path beliefs may not be fully determined by Bayes' rule and equilibrium strategies — Bayes rule doesn't tell us how to update after probability 0 events.
  - Subgame perfection still implies that players ignore threats involving strategies that are strictly dominated beginning at any information set, even off the equilibrium path.
- The definition of a PBE in a signaling game is: $\{m^*(\theta), a^*(m), \mu(\theta|m)\}$ such that:

  1. Sender:

$$m^*(\theta) = \arg\max_{m \in M} \pi_S(m, a^*(m), \theta) \tag{7}$$

- Sender chooses a payoff-maximizing message, m^*, given her type θ and given that Respondent will choose his optimal response $a^*(m)$ to that message.

2. Respondent:

$$a^*(m) = \arg\max_{a \in A} E_\mu \pi_R(a, m, \theta) = \sum_{\theta \in \Theta} \mu(\theta|m)\pi_R(a, m, \theta) \tag{8}$$

   - Respondent chooses a payoff-maximizing action $a^*(m)$, given his posterior beliefs $\mu(\theta|m)$ after receiving the message.

3. Beliefs $\mu(\theta|m)$ are consistent with Bayes' rule, given the prior distribution of θ and Sender's strategy (7).

- Sender chooses a message that maximizes her payoff when the Responder chooses an action (given that message) that maximizes his payoff.

- So to check for PBE, we have to make sure that the hypothesized message actually is a best-response for the Sender.

Pooling Equilibria

Recall the Responder's best-responses in the example:

$$a^*(l) = u \quad \forall p \in [0, 1] \tag{9}$$

$$a^*(r) = \begin{cases} d & \text{if } q \leq \frac{2}{3} \\ u & \text{if } q \geq \frac{2}{3} \end{cases} \tag{10}$$

Begin by guessing that the Sender uses the pooling strategy r.

1. Sender strategy is $s_1 = (r|\theta_H, r|\theta_L)$.

2. Then $\mu(\theta|m)$ is equal to the priors, thus $q = 0.5$.

3. Given this: $a^*(r) = d$.

4. Given this response, the two sender types have payoffs from message r:

$$\begin{aligned} \pi_S(s_1, a^*(r)|\theta_H) &= 0 \\ \pi_S(s_1, a^*(r)|\theta_L) &= 1 \end{aligned} \tag{11}$$

5. Given these payoffs we can ask if s_1 is a best-response, or would either of the two types benefit from deviating to a different strategy?

6. Suppose that the θ_H type instead sent message l.

   - Recall again that $a^*(l) = u \quad \forall p \in [0, 1]$.
   - Thus, the payoff from a θ_H type sending l is $1 > 0$, and deviating from s_1 is indeed beneficial.
   - That is, s_3 or s_4 dominates s_1, so pooling on r (i.e., using s_1) can not happen in PBE.

Let's now look for pooling on l.

1. This is the sender strategy: $s_4 = (l|\theta_H, l|\theta_L)$.

2. Again, all types choose the same message so $\mu(\theta|m)$ is equal to the priors, thus $p = 0.5$.

3. Since $m = r$ has probability 0 in this hypothesized equilibrium, there is no constraint on q. It is a free parameter.

4. Responder plays $a^*(l) = u \quad \forall p \in [0, 1]$.

5. so Sender's payoffs in this case will be:

$$\begin{aligned} \pi_S(s_4, a^*(l)|\theta_H) &= 1 \\ \pi_S(s_4, a^*(l)|\theta_L) &= 2 \end{aligned}$$

6. But, can Sender benefit by deviating?

    - To figure this out, we need to see how Respondent reacts to the other message.
    - If Respondent chooses $u|r$ (strategies r_1 or r_2), which is a best-response for all $q \geq \frac{2}{3}$, then a θ_H sender would earn:

    $$\pi_S(r, u|r, \theta_H) = 2 > 1,$$

    so the deviation would benefit her if $q \geq \frac{2}{3}$.
    - If Respondent instead chooses $d|r$ (strategies r_3 or r_4), which is a best-response for all $q \leq \frac{2}{3}$, then a θ_H sender would earn:

    $$\pi_S(r, d|r, \theta_H) = 0 < 1,$$

    not beneficial.
    - For the θ_L types if $d|r$ (strategies r_3 or r_4), which is a best-response for all $q \leq \frac{2}{3}$, then a θ_L sender would earn:

    $$\pi_S(r, d|r, \theta_L) = 1 < 2,$$

    again not beneficial.
    - Therefore, in PBE with pooling on l (that is, where Sender plays s_4) we must have $a^*(r) = d$, which requires $q \leq \frac{2}{3}$.
        - That is, for this PBE to work, Responder must play strategy $r_3 = (d|r, u|l)$, which requires beliefs $q \leq \frac{2}{3}$.

- Thus, we have a pooling PBE:

$$\text{Pooling PBE} = [s_4 = (l|\theta_H, l|\theta_L), r_3 = (d|r, u|l), \mu] \tag{12}$$

where beliefs μ are $p = 0.5$ and any $q \leq \frac{2}{3}$.

- Note that only the l-information set is on the equilibrium path, so the posterior probability $p = 0.5$ is equal to the prior, while q can take any value.

Separating Equilibria

The big difference here is that beliefs are zero or one since there are only two signals and two types.

Look first at $s_3 = (l|\theta_H, r|\theta_L)$. Can this be part of a PBE profile?

1. If sender plays s_3 then both information sets are on the equilibrium path.

2. Beliefs are determined by Bayes' rule and Sender's strategy s_3, thus $p = 1, q = 0$.

3. Given these beliefs:

$$
\begin{aligned}
a^*(l) &= u \\
a^*(r) &= d
\end{aligned}
$$

4. Now Responder's best-response is: $r_3 = (d|r, u|l)$ and Senders payoff is:

$$
\begin{aligned}
\pi_S(s_3, a^*(l), \theta_H) &= 1 \\
\pi_S(s_3, a^*(r), \theta_L) &= 1
\end{aligned}
$$

5. Now we can check: given that she faces r_3, is s_3 a best-response for Sender?

6. No, not for the θ_L type. If she deviates and chooses l then Respondent plays r_3, $a^*(l) = u$ and

$$\pi_S(a^*(l), \theta_L) = 2 > 1$$

7. Therefore s_3 is not optimal and there is no separating equilibrium in which $s_3 = (l|\theta_H, r|\theta_L)$ is played.

Finally, let's look for a separating equilibrium using $s_2 = (r|\theta_H, l|\theta_L)$.

1. If s_2 is played in equilibrium then the respondent's beliefs must be: $p = 0$, $q = 1$.

2. Therefore, the respondents best-response is:

$$
\begin{aligned}
a^*(l) &= u \\
a^*(r) &= u
\end{aligned}
$$

3. which implies strategy $r_1 = (u|r, u|l)$.

4. Payoffs for the two sender types are then:

$$
\begin{aligned}
\pi_S(r|\theta_H, a^*(r|q=1), \theta_H) &= 2 \\
\pi_S(l|\theta_L, a^*(l|p=0), \theta_L) &= 2
\end{aligned}
$$

5. If the θ_H type were to deviate and play l,then $a^*(l) = u$ and $\pi_S(l|\theta_H, a^*(r|q = 1), \theta_H) = 1 < 2$, thus there is no incentive for a θ_H type to deviate.

6. If the θ_L type were to deviate and play r, then $a^*(r) = u$ and $\pi_S(r|\theta_L, a^*(r|q = 1), \theta_L) = 1 < 2$, thus there is no incentive for a θ_L type to deviate.

Thus, we have a separating PBE:

$$[s_2 = (r|\theta_H, l|\theta_L), r_1 = (u|r, u|l), p = 0, q = 1] \tag{13}$$

To conclude, this example game has two PBE: pooling on l with Responder beliefs $q \leq \frac{2}{3}$,[1] and separating with the high type Sender choosing r with Responder beliefs $p = 0, q = 1$.

[1]Properly speaking, this is a continuum of pooling PBE that differ only in the value of q, but informally game theorists refer to it as a single PBE.

6 Screening Models

- Screening: The uniformed party creates a menu of options.
  - In a separating equilibrium the informed party's choice of action reveals its type.
  - In a pooling equilibrium all types choose the same action and the screening is unsuccessful.
- Denote:
  - U- uninformed party
  - I- informed party
  - N- nature
  - Timing of a signaling model: N, I, U.
  - Timing of a screening model: U, N, I; or N, U, I with U not observing nature's realization.

Screening Lemons

- Can we reformulate Akerlof's Lemons model as a screening model?
- Yes, via a deferred payment plan.
  - Recall that $\theta_1 = 2$ was the lower bound on the value for a buyer.
  - Pay $\$2,000$ now and put x in escrow.
  - Whatever is left from the escrow account after paying repairs goes to the seller.

- Buyer offers the contract:

$$\pi_I = 2,000 + (x - repairs)$$
$$\pi_U = \theta - 2,000 - x$$

  Where the seller is the informed party (π_I) and the buyer is the uninformed party (π_U).

- Nature chooses the repairs, assuming repairs are less than x.

- Ex: $\theta_i = 2,500$, $x = 500$, and repairs $= 200$. $\pi_I = 2,000 + (500 - 200) = 2,300$, and $\pi_U = 2,500 - 2,000 - 500 = 0$.
  - Example might need some further work.
  - Note that all the surplus goes to the seller in this example since the contract awards the entire amount remaining in the escrow account to the seller.
  - If this were shared then the payoff for the buyer would be positive.
  - Also, note that the sellers reservation value is: $r(\theta) = \theta - 0.1$, so it is very possible that the seller could be worse off by selling the car if there are a lot of repairs.
  - Since the seller incurs the risk, they receive the entire escrow account and given the probability distribution for the repair cost we require that the contract at least satisfies the sellers certainty equivalent.

7 Insurance model

- For simplicity, let there be two types of customers, high or low risk: $\Theta = \{hi, lo\}$.
- The consumer pays a premium, that is the price for insurance, which is a "bad" in the sense that they are worse off with a higher premium, so we can think of a reduction in the premium as a "good."

- The consumer would prefer a low deductible, so the deductible is a "bad."
  - We will draw the indifference maps with a good on the vertical axis (reduction in the premium) and a bad on the horizontal axis (amount of the deductible).
  - This means that we have positively sloped indifference curves.
  - We will see the same sort of ICs in the next chapter for the standard moral hazard (hidden action) model, where the workers wage is a good (vertical axis) and effort has a cost and is therefore a bad.
- We can draw the indifference maps for the two types, with the high risk types having steeper indifference curves and utility is increasing to the northwest (rather than the northeast with two goods on the axes), $I_L^1 < I_L^2 < I_L^3 < I_L^4$, and $I_H^1 < I_H^2 < I_H^3 < I_H^4$.

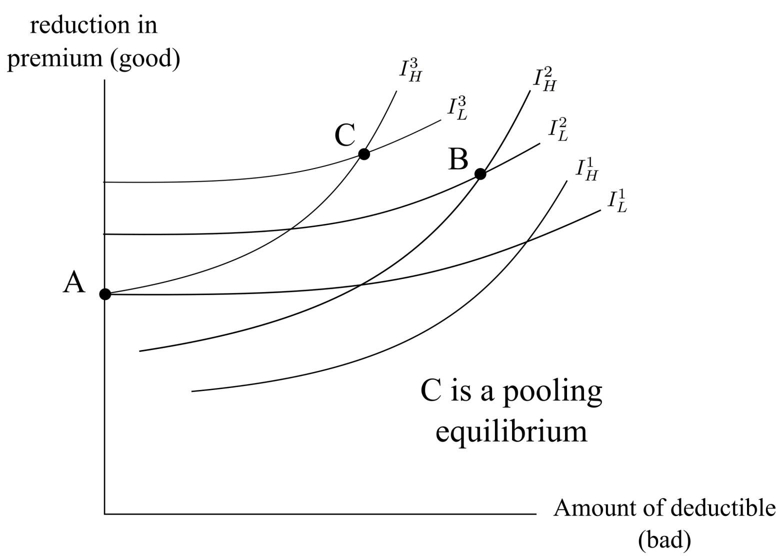

Figure 2: Indifference curves for both types, and contract menu A, B, C.

- The high risk types have steep indifference curves compared to the low risk types. Why?

  - The premium is paid with certainty while the deductible is only relevant if there is an accident.
  - Thus, a low risk type is less likely to use the deductible and therefore they will be more "tolerant" of the bad.

- We can look at point A at the intersection of the vertical axis, I_L^1 and I_H^4 as being a high premium and no deductible.

- Point B in the interior has a lower premium and a higher deductible at the intersection of I_L^2 and I_H^3.

- In this scenario, the high risk types will choose point A and the low risk types choose point B.
  - Thus, we have a separating equilibrium for our screening model.
  - The informed party (consumer) reveals their type, and hence their private information, by their choice of contract.

- Suppose we offered a menu with A, B, and C. Now all low risk types will choose C and high risk types are indifferent between A and C. This implies partial pooling on C, with a Bayesian posterior of one that an individual is high risk if A is chosen. The posteriors will be greater than the priors if C is chosen, but less than one.

- We have self-selection.
  - It is harder to get pooling in a screening model than it is in a signaling model.
  - It might be the case that the signal is often costly to send (ex: obtaining a college degree has a substantial cost) while setting up a menu of contracts is relatively costless.

- Formally, we have two constraints.

1. Incentive constraint:

$$\begin{aligned} E(U_{\theta_H}(A)) &\geq E(U_{\theta_H}(B)) \\ E(U_{\theta_L}(B)) &\geq E(U_{\theta_L}(A)) \end{aligned} \tag{14}$$

   - Thus, the high risk types choose A and the low risk types choose B.
   - There is an incentive for the informed party to truthfully reveal their type.
   - The market failure due to asymmetric information is overcome.

2. Participation constraint:

$$\begin{aligned} E(U_{\theta_H}(A)) &\geq r(\theta_H) \\ E(U_{\theta_L}(B)) &\geq r(\theta_L) \end{aligned} \tag{15}$$

   where $r(\theta_i)$ is the reservation utility for type θ_i.

   - The uninformed party creates a menu of options that maximize profit subject to these two constraints.

- Another example: [I deleted the first one][[Matt, let's discuss this last sentence...it doesn't quite make sense to me.]][Sorry. This was from a paper by a colleague here at UWM and familiar to some of our students.]Labor: Workers ability is private information.

  1. – Effort is more costly for the low ability types.
     – A firm will offer a contract with high wages, but high effort required, and another contract targeted at low ability types with lower required effort and lower wages.

8 Further Reading

Spence, M. (1973). "Job Market Signaling," The Quarterly Journal of Economics, 87(3): 355-74.

Wilson, C. (1980). “The Nature of Equilibrium in Markets with Adverse Selection,” The Bell Journal of Economics, 11(1): 108-30.