

ADVANCED MICROECONOMICS: LECTURE NOTE 14

Instructor: Xiang Sun

2024 Spring

1 Motivating example 1

1. 项目经理手头有一些项目（两种可能的结果：成功，失败），想说服领导批准。¹

- 假设每个项目的成功概率为 30%，这是领导和项目经理的共同知识。假设领导只在成功率大于 50% 的情况下会批准。
- 如果项目经理发起了详细的调查研究，详详细细地把所有的项目都解释给领导听，领导获得了完全的信息，也只会批准 30% 的项目，另外 70% 的项目就会被拒绝。
- 作为项目经理，自然希望项目批准的越多越好。那么用什么策略能够增加领导批准的概率呢？信息发掘到极点，也就 30% 的可能性，似乎面对理性的领导，说服是多余而无效的。

每个项目可以被视为有两种类型：“成功”和“失败”；前者的概率是 30%。

2. 项目经理可以承诺如下“机制”：

- 首先，项目经理可以这么说：“保险起见，这些项目我进行了预审核。假如项目最终能成功，那么该项目一定在我的推荐范围之内，领导永远不会因为听我的推荐而错失能够成功的项目。”

$$\text{Prob}(\text{推荐} \mid \text{成功}) = 1$$

- 接下来项目经理又说了：“至于我推荐的呢，我不敢保证所有的都能成功，但是经过详细的调查，我保证我推荐的都是成功概率大的。假如这个项目会失败，那么我大约有 $\frac{4}{7}$ 的概率会不给您推荐， $\frac{3}{7}$ 的概率会给您推荐。”

$$\text{Prob}(\text{推荐} \mid \text{失败}) = \frac{3}{7}$$

领导一想，这个也是合理的，毕竟你可能也无法精确地讲解出什么一定成功，什么一定失败；反正 $\frac{4}{7} > \frac{3}{7}$ ，你的推荐似乎还是很有意义的。

项目经理有两个行为：“推荐”、“不推荐”。他事先确定机制或者信息规则 $\text{Prob}(\text{推荐} \mid \text{成功})$ 和 $\text{Prob}(\text{推荐} \mid \text{失败})$ 。

3. 如果领导接受上面所有的设定并且觉得没什么，那么恭喜项目经理——领导被成功地说服了，领导最终会批准 60% 的项目！比预期整整多了一倍！

¹本例选自<https://zhuanlan.zhihu.com/p/25275142>。

- (a) 任意一个项目，推荐的概率 = 能成功而推荐的概率 × 能成功概率 + 会失败而推荐的概率 × 会失败的概率 = $1 \times 30\% + \frac{3}{7} \times 70\% = 60\%$ ：

$$\begin{aligned}\text{Prob}(\text{推荐}) &= \text{Prob}(\text{推荐} | \text{成功}) \times \text{Prob}(\text{成功}) + \text{Prob}(\text{推荐} | \text{失败}) \times \text{Prob}(\text{失败}) \\ &= 1 \times 30\% + \frac{3}{7} \times 70\% = 60\%.\end{aligned}$$

- (b) 而一旦项目经理推荐了，领导就会听，因为：推荐而最后成功的概率 = 能成功而推荐的概率 × 能成功的概率 ÷ 推荐的概率 = $1 \times 30\% \div 60\% = 50\%$ 。

$$\text{Prob}(\text{成功} | \text{推荐}) = \frac{\text{Prob}(\text{推荐} | \text{成功}) \times \text{Prob}(\text{成功})}{\text{Prob}(\text{推荐})} = \frac{1 \times 30\%}{60\%} = 50\%.$$

- (c) 也就是说现在只要项目经理推荐了，能成功的概率就刚刚好等于 50%，现实中可以稍微的比 50% 多一点，那么领导就一定会批准了。
- (d) 所以在这种情况下，项目经理以略微低于 60% 的概率推荐项目，而领导会听从项目经理的所有推荐，因为只要项目经理推荐了，那成功概率就略大于 50%。

4. 项目经理的推荐方式某种意义上是无偏的。从事前来看，领导觉得项目会“成功”的概率为：

$$\text{Prob}(\text{不推荐})\text{Prob}(\text{成功} | \text{不推荐}) + \text{Prob}(\text{推荐})\text{Prob}(\text{成功} | \text{推荐}) = 0.4 \times 0 + 0.6 \times \frac{1}{2} = 0.3$$

这与初始给定的先验概率一致。

因此，在领导是“理性地”批准了 60% 的项目，同时他知道其实只有 30% 的项目能成功。这个过程中，所有人都是完全理性的。

5. 为什么是这种推荐方式？

- (a) 首先考察 $\text{Prob}(\text{成功} | \text{推荐})$ 。

$$\begin{aligned}\text{Prob}(\text{成功} | \text{推荐}) &= \frac{\text{Prob}(\text{推荐} | \text{成功}) \times \text{Prob}(\text{成功})}{\text{Prob}(\text{推荐} | \text{成功}) \times \text{Prob}(\text{成功}) + \text{Prob}(\text{推荐} | \text{失败}) \times \text{Prob}(\text{失败})} \\ &= \frac{3 \text{Prob}(\text{推荐} | \text{成功})}{3 \text{Prob}(\text{推荐} | \text{成功}) + 7 \text{Prob}(\text{推荐} | \text{失败})}\end{aligned}$$

注意到 $\text{Prob}(\text{成功} | \text{推荐})$ 关于 $\text{Prob}(\text{推荐} | \text{成功})$ 严格递增。更高的 $\text{Prob}(\text{推荐} | \text{成功})$ 会从两个方面使得项目经理受益：

- “推荐”的无条件概率提高了。
- “推荐”更有说服力——它增加了领导的后验猜测 $\text{Prob}(\text{成功} | \text{推荐})$ 。

因此，项目经理的最优选择就是让 $\text{Prob}(\text{推荐} | \text{成功})$ 最大，即设为 1。

- (b) 最优的规则将保证：当项目经理“推荐”时，领导在“批准”与“不批准”之间无差异。

- 一方面，如果此时领导严格偏好“批准”，那么说明 $\text{Prob}(\text{成功} | \text{推荐}) > 0.5$ ，项目经理用力过度：项目经理可以通过增加 $\text{Prob}(\text{推荐} | \text{失败})$ ，在保证 $\text{Prob}(\text{成功} | \text{推荐}) \geq 0.5$ 的前提下，增加“推荐”的无条件概率。

- 另一方面，如果此时领导严格偏好“不批准”，那么说明 $\text{Prob}(\text{成功} | \text{推荐}) < 0.5$ ，项目经理竹篮打水一场空。

(c) 于是

$$0.5 = \text{Prob}(\text{成功} | \text{推荐}) = \frac{3 \text{Prob}(\text{推荐} | \text{成功})}{3 \text{Prob}(\text{推荐} | \text{成功}) + 7 \text{Prob}(\text{推荐} | \text{失败})} = \frac{3}{3 + 7 \text{Prob}(\text{推荐} | \text{失败})}$$

所以 $\text{Prob}(\text{推荐} | \text{失败}) = \frac{3}{7}$ 。

- 一方面，只要领导的后验猜测 $\text{Prob}(\text{成功} | \text{推荐}) = 0.5$ ，就没有必要再增加了。
- 另一方面，如果领导的后验猜测 $\text{Prob}(\text{成功} | \text{推荐}) < 0.5$ ，那么“推荐”。

6. 关于领导的更细致分析，可以帮助我们发现其中的直觉。

(a) 我们先为领导构造一个合适的效用函数，便于后面计算：可以验证，当领导的猜测为 $\text{Prob}(\text{成功}) =$

	成功	失败
推荐	1	-1
不推荐	0	0

0.5 时，他在“批准”与“不批准”之间无差异。

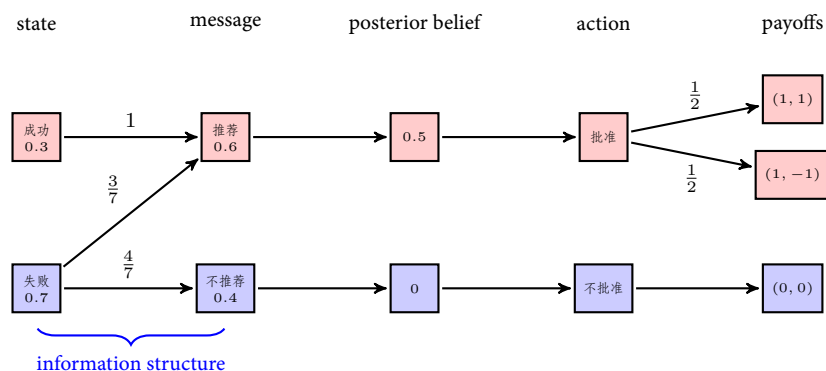
(b) 如果领导完全不考虑项目经理的信息，由于初始猜测为 0.3，因此，他将“不批准”，效用为零。

(c) 如果领导考虑经理的信息，

- 情况 1：他有 60% 的概率收到“推荐”，这时他的后验猜测调整为 $\text{Prob}(\text{成功} | \text{推荐}) = 0.5$ ，他将“批准”，获得期望效用 $0.5 \times 1 + 0.5 \times (-1) = 0$ 。
- 情况 2：他有 40% 的概率收到“不推荐”，这时他的后验猜测调整为 $\text{Prob}(\text{成功} | \text{不推荐}) = 0$ ，他将“不批准”，获得期望效用 0。
- 因此，他的事前期望效用为零。

不考虑项目经理的信息时，领导永远会“不批准”。考虑项目经理的信息时，如果项目是“成功”，由于项目经理的“推荐”，领导将不会错过，从而获得正收益；另一方面，他会因为“推荐”而“批准”太多项目，其中包括（一半的）“失败”项目，这会有负收益。两者正如情况 1 中恰好抵消。

7. 总结



8. 如果先验估计是 0.6，结论又如何？

2 Motivating example 2

9. Consider an example of a prosecutor (sender) trying to convince a judge (receiver) that a defendant is guilty.

- When the defendant is indeed guilty, revealing the facts of the case will tend to help the prosecutor's case.
- When the defendant is innocent, revealing facts will tend to hurt the prosecutor's case.

Can the prosecutor “structure” his arguments, selection of evidence, etc. so as to increase the probability of conviction by a rational judge on average?

10. There are two states of the world: the defendant is either **guilty** or **innocent**. The judge must choose one of two actions: to **acquit** or **convict**.

- The judge gets utility 1 for choosing the just action (convict when guilty and acquit when innocent) and utility 0 for choosing the unjust action (convict when innocent and acquit when guilty).
- The prosecutor gets utility 1 if the judge convicts and utility 0 if the judge acquits, regardless of the state.
- The prosecutor and the judge share a prior belief $\mu_0(\text{guilty}) = 0.3$.

11. The prosecutor conducts an investigation and is required by law to **report its full outcome**.

We can think of the choice of the investigation as consisting of the decisions on whom to subpoena, what forensic tests to conduct, what questions to ask an expert witness, etc.

We formalize an **investigation** as distributions $q(\cdot | \text{guilty})$ and $q(\cdot | \text{innocent})$ on some set of signal realizations. The prosecutor chooses q and must honestly report the signal realization to the judge.

12. If the prosecutor chooses a fully informative investigation, one that leaves no uncertainty about the state, the judge convicts 30 percent of the time.

13. However, the prosecutor can do better. For example, he can choose the following binary signal $\{i, g\}$ such that

$$q(g | \text{guilty}) = 1 \text{ and } q(g | \text{innocent}) = \frac{3}{7}.$$

(a) The first signal g leads to a posterior

$$\mu_g = \mu(\cdot | g) = \frac{\mu_0(\cdot) \cdot q(g | \cdot)}{\mu_0(\text{guilty}) \cdot q(g | \text{guilty}) + \mu_0(\text{innocent}) \cdot q(g | \text{innocent})} = \frac{1}{2}\delta_{\text{innocent}} + \frac{1}{2}\delta_{\text{guilty}}.$$

That is, the sender says “this person is guilty enough to convict”: After observing g and knowing the posterior μ_g , her optimal action is convict, and the expected utility is $\frac{1}{2}$ (By default, we assume that optimal action is convict once convict and acquit are indifferent).

Notice that the judge has probability $\frac{1}{2}$ to choose an unjust action.

(b) The second signal i leads to a posterior

$$\mu_i = \mu(\cdot | i) = \frac{\mu_0(\cdot) \cdot q(i | \cdot)}{\mu_0(\text{guilty}) \cdot q(i | \text{guilty}) + \mu_0(\text{innocent}) \cdot q(i | \text{innocent})} = \delta_{\text{innocent}}.$$

That is, the sender says “this person is innocent”: After observing i and knowing the posterior μ_i , her optimal action is acquit, and the expected utility is 1.

(c) Moreover, the probability of signal g is

$$\text{Prob}(g) = \mu_0(\text{innocent}) \cdot q(g \mid \text{innocent}) + \mu_0(\text{guilty}) \cdot q(g \mid \text{guilty}) = 0.7 \times \frac{3}{7} + 0.3 \times 1 = 0.6,$$

and the probability of signal i is $\text{Prob}(i) = 0.4$.

This leads the judge to convict with (*ex ante* level) probability 60%. The judge knows 70% of defendants are innocent, yet she convicts 60% of them!

(d) Furthermore, she does so even though she is fully aware that the investigation was designed to maximize the probability of conviction, or the investigation is unbiased.

The expected utility of the judge is $0.6 \times \frac{1}{2} + 0.4 \times 1 = 0.7$. When there is no persuasion, the optimal action of the judge is acquit, and the expected utility is also 0.7. That is, the optimal expected utilities of the judge are the same. This equivalence can be formally stated as the following statement:

$$\mu_0(\theta) = \text{Prob}(g) \cdot \mu(\theta \mid g) + \text{Prob}(i) \cdot \mu(\theta \mid i)$$

for $\theta = \text{guilty}$ or innocent .

3 A model of Bayesian persuasion

14. There is a sender and a receiver. The state space Θ is finite and the action space is compact. The sender and the receiver share the (full support) common prior $\mu_0(\cdot)$ on Θ .
15. Receiver has a continuous utility function $u_r(a, \theta)$, and sender has a continuous utility function $u_s(a, \theta)$.
16. The special feature of the model is that the sender's strategy is to pick an **information structure** (or **information disclosure rule**),
 - a finite set M of messages,
 - a mapping $q: \Theta \rightarrow \Delta(M)$.

Here, $q(m \mid \theta)$ describes the probability that the receiver hears the message m when θ is the true state.

Note that the mapping q can be induced by a joint distribution on states and messages.

事先承诺一套信息规则。

17. The persuasion game works as follows.
 - (1) The sender selects the information structure (M, q) .
 - (2) Nature picks the true state according to prior μ_0 .
 - (3) The information structure (M, q) generates a message m .
 - (4) Given the information structure, the receiver selects an action rule (what to do following any message).

The sender is potentially informed but can choose to commit to public information acquisition and disclosure.

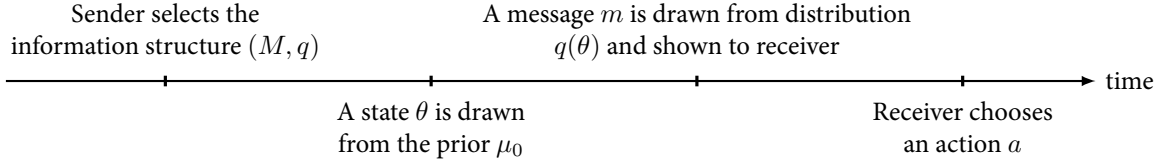


Figure 1: Timing

18. In equilibrium:

- The receiver's action rule is a best response to the information structure (for each realized message).

After observing the sender's choice of information structure (M, q) and a message realization $m \in M$ (drawn by nature), the receiver **forms the posterior** $\mu_m(\cdot) \in \Delta(\Theta)$ **via Bayes' rule**:

$$\mu_m(\theta) = \frac{\mu_0(\theta) \cdot q(m | \theta)}{\sum_{\theta' \in \Theta} \mu_0(\theta') \cdot q(m | \theta')},$$

and **takes an action** $\hat{a}(\mu_m) \in A$ to maximize her expected utility

$$\mathbb{E}_{\mu_m}[u_r(a, \theta)] = \sum_{\theta \in \Theta} \mu_m(\theta) \cdot u_r(a, \theta).$$

If there are multiple best actions, assume that the receiver picks one of the optimizers that maximizes the sender's expected utility $\mathbb{E}_{\mu_m}[u_s(a, \theta)] = \sum_{\theta \in \Theta} \mu_m(\theta) \cdot u_s(a, \theta)$.

- The sender selects the information structure (M, q) to maximize her expected utility given the receiver's response.

Given an information structure (M, q) , let

$$P(m) = \sum_{\theta \in \Theta} q(m | \theta) \cdot \mu_0(\theta)$$

be the probability that message m is heard.

The sender's *ex ante* expected utility is

$$U_s(M, q) = \mathbb{E}_P \left[\mathbb{E}_{\mu_m} \left[u_s(\hat{a}(\mu_m), \theta) \right] \right] = \sum_{m \in M} P(m) \left[\sum_{\theta \in \Theta} \mu_m(\theta) \cdot u_s(\hat{a}(\mu_m), \theta) \right].$$

Given the receiver's action rule, the sender **chooses the information structure** (M, q) **to maximize his expected utility**.

19. Key assumption: The sender **cannot distort or conceal information** once the message is realized.

20. We simplify the analysis further by noting that, without loss of generality, we can **restrict our attention to a particular class of information structures**. Say that an information structure is straightforward if $M \subseteq A$ and the receiver's equilibrium action equals the message realization.

In other words, a straightforward information structure produces a "recommended action" and the receiver always follows the recommendation.

21. Lemma: The following are equivalent:

- (i) There exists an information structure with the expected utility u_s^* .

(ii) There exists a straightforward information structure with expected utility u_s^* .

这个结果与机制设计中的显示原理类似，可以让我们将注意力放在“straightforward information structure”上。

Proof. (1) By definition, (ii) implies (i).

(2) Given an information structure (M, q) with the expected utility u_s^* , let

$$M^a = \{m \in M \mid \hat{a}(\mu_m) = a\} \text{ for each } a \in A.$$

(3) Consider an information structure (M', q') with $M' = A$ and

$$q'(a \mid \theta) = \sum_{m \in M^a} q(m \mid \theta).$$

(4) For each $a \in A$, we have

$$\begin{aligned} \sum_{\theta \in \Theta} \mu'_a(\theta) \cdot u_r(a, \theta) &= \sum_{\theta \in \Theta} \frac{\mu_0(\theta) \cdot \sum_{m \in M^a} q(m \mid \theta)}{\sum_{\theta' \in \Theta} \mu_0(\theta') \cdot \sum_{m' \in M^a} q(m' \mid \theta')} \cdot u_r(a, \theta) \\ &= \frac{1}{\sum_{\theta' \in \Theta} \mu_0(\theta') \cdot \sum_{m' \in M^a} q(m' \mid \theta')} \sum_{\theta \in \Theta} \mu_0(\theta) \sum_{m \in M^a} q(m \mid \theta) \cdot u_r(a, \theta) \\ &= \frac{1}{\sum_{\theta' \in \Theta} \mu_0(\theta') \cdot \sum_{m' \in M^a} q(m' \mid \theta')} \sum_{\theta \in \Theta} \mu_0(\theta) \sum_{m \in M^a} q(m \mid \theta) \cdot u_r(\hat{a}(\mu_m), \theta) \\ &\geq \frac{1}{\sum_{\theta' \in \Theta} \mu_0(\theta') \cdot \sum_{m' \in M^a} q(m' \mid \theta')} \sum_{\theta \in \Theta} \mu_0(\theta) \sum_{m \in M^a} q(m \mid \theta) \cdot u_r(a', \theta) \\ &= \sum_{\theta \in \Theta} \mu'_a(\theta) \cdot u_r(a', \theta) \end{aligned}$$

for each $a' \in A$. That is, a is an optimal response for the receiver to the realization a from q' .

(5) Sender's expected utility under the new information structure is

$$\begin{aligned} U_s(M', q') &= \sum_{a \in A} P(a) \left[\sum_{\theta \in \Theta} \mu'_a(\theta) \cdot u_s(\hat{a}(\mu'_a), \theta) \right] \\ &= \sum_{a \in A} P(a) \left[\sum_{\theta \in \Theta} \mu'_a(\theta) \cdot u_s(a, \theta) \right] \\ &= \sum_{a \in A} \left(\sum_{\theta' \in \Theta} \sum_{m \in M^a} q(m \mid \theta') \cdot \mu_0(\theta') \right) \left[\sum_{\theta \in \Theta} \frac{\mu_0(\theta) \sum_{m \in M^a} q(m \mid \theta)}{\sum_{\theta'' \in \Theta} \mu_0(\theta'') \sum_{m \in M^a} q(m \mid \theta'')} u_s(a, \theta) \right] \\ &= \sum_{a \in A} \sum_{\theta \in \Theta} \mu_0(\theta) \sum_{m \in M^a} q(m \mid \theta) u_s(\hat{a}(\mu_m), \theta) \\ &= \sum_{m \in M} P(m) \left[\sum_{\theta \in \Theta} \mu_m(\theta) \cdot u_s(\hat{a}(\mu_m), \theta) \right] = U_s(M, q) \end{aligned}$$

□

4 Simplified problem

22. Each message m induces a posterior belief $\mu_m(\cdot)$. Thus, an information structure (M, q) naturally induces a distribution over posterior beliefs: for each posterior μ ,

$$\tau(\mu) = \text{Prob}(\mu) = \sum_{m: \mu_m = \mu} \text{Prob}(m) = \sum_{m: \mu_m = \mu} \sum_{\theta' \in \Theta} q(m | \theta') \cdot \mu_0(\theta').$$

类比例子，后验猜测完全由信号决定。

- 信号“推荐”对应着后验猜测 μ_r “ $\frac{1}{2}$ 成功、 $\frac{1}{2}$ 失败”；
- 信号“不推荐”对应着后验猜测 μ_n “100% 失败”。

两个后验猜测不一样，各自出现的概率与对应的信号的概率一致：后验猜测 μ_r 的概率 60%，而后验猜测 μ_n 的概率 40%——这就是“所有后验猜测集合”上的一个分布，即给每个后验猜测赋予出现的概率。

信号	后验猜测	概率
“推荐”	μ_r	0.6
“不推荐”	μ_n	0.4

Conversely, a distribution over posterior beliefs $\tau \in \Delta(\Delta(\Theta))$ is induced by an information structure (M, q) if

$$\begin{aligned} \mu_m(\theta) &= \frac{\mu_0(\theta) \cdot q(m | \theta)}{\sum_{\theta' \in \Theta} \mu_0(\theta') \cdot q(m | \theta')} \text{ for all } \theta \text{ and } m, \\ \text{support}(\tau) &= \{\mu_m\}_{m \in M}, \\ \tau(\mu) &= \sum_{m: \mu_m = \mu} \sum_{\theta' \in \Theta} q(m | \theta') \cdot \mu_0(\theta') \text{ for all } \mu. \end{aligned}$$

23. When both sender and receiver hold some posterior μ , the sender's expected utility is equal to

$$U_s(\mu) = \mathbf{E}_\mu \left[u_s(\hat{a}(\mu), \theta) \right] = \sum_{\theta \in \Theta} \mu(\theta) \cdot u_s(\hat{a}(\mu), \theta).$$

如果双方拥有相同的后验猜测 μ ，那么基于相同的推断过程，receiver 将选择 $\hat{a}(\mu)$ ，同时 sender 也知道 receiver 将选择 $\hat{a}(\mu)$ 。这时，sender 效用中唯一的不确定因素是 θ ——他将使用给定的后验猜测 μ 来估计。

24. Since sender's and receiver's beliefs coincide, the sender's utility from any distribution of posteriors τ (induced by some information structure) is

$$U_s(\tau) = \mathbf{E}_\tau U_s(\mu) = \mathbf{E}_\tau \left[\mathbf{E}_\mu \left[u_s(\hat{a}(\mu), \theta) \right] \right] = \sum_{\mu \in \text{support}(\tau)} \tau(\mu) \left[\sum_{\theta \in \Theta} \mu(\theta) \cdot u_s(\hat{a}(\mu), \theta) \right].$$

更进一步，基于“后验猜测”上的猜测 τ ，sender 将评估自己的期望收益：每个后验猜测 μ 出现的概率是 $\tau(\mu)$ ，每个 μ 对应着效用 $U_s(\mu)$ 。

25. A distribution of posteriors τ is **Bayesian plausible** if the expected posterior probability equals the prior:

$$\mathbf{E}_\tau \mu = \sum_{\mu \in \text{support}(\tau)} \tau(\mu) \cdot \mu = \mu_0.$$

本质上，这就是无偏条件。回忆例子，领导觉得项目会成功的概率为：

$$\text{Prob}(\text{不推荐}) \text{Prob}(\text{成功} \mid \text{不推荐}) + \text{Prob}(\text{推荐}) \text{Prob}(\text{成功} \mid \text{推荐}) = 0.4 \times 0 + 0.6 \times \frac{1}{2} = 0.3$$

这与初始给定的先验概率一致。

26. Clearly, if τ is induced by an information structure (M, q) , then τ is Bayesian plausible.

For each θ ,

$$\begin{aligned} \sum_{\mu \in \text{support}(\tau)} \tau(\mu) \cdot \mu(\theta) &= \sum_{\mu \in \text{support}(\tau)} \sum_{m: \mu_m = \mu} \sum_{\theta' \in \Theta} q(m \mid \theta') \cdot \mu_0(\theta') \cdot \mu(\theta) \\ &= \sum_{\mu \in \text{support}(\tau)} \sum_{m: \mu_m = \mu} q(m \mid \theta) \cdot \mu_0(\theta) = \mu_0(\theta) \end{aligned}$$

27. Proposition: The following are equivalent:

- (i) There exists an information structure (M, q) such that the expected utility $U_s(M, q)$ is u_s^* .
- (ii) There exists a Bayesian plausible distribution of posteriors τ such that the expected utility $U_s(\tau)$ is u_s^* .

这个结果将问题进一步化简：从寻找 information structure，转化为寻找 distribution of posteriors。

Proof. (1) Given an information structure (M, q) , let τ be the distribution of posteriors induced by (M, q) . Then

$$\begin{aligned} U_s(\tau) &= \sum_{\mu \in \text{support}(\tau)} \tau(\mu) \left[\sum_{\theta \in \Theta} \mu(\theta) \cdot u_s(\hat{a}(\mu), \theta) \right] \\ &= \sum_{\mu \in \text{support}(\tau)} \left[\sum_{m: \mu_m = \mu} \sum_{\theta' \in \Theta} q(m \mid \theta') \cdot \mu_0(\theta') \right] \cdot \left[\sum_{\theta \in \Theta} \mu(\theta) \cdot u_s(\hat{a}(\mu), \theta) \right] \\ &= \sum_{\mu \in \text{support}(\tau)} \sum_{\theta \in \Theta} \sum_{m: \mu_m = \mu} \sum_{\theta' \in \Theta} q(m \mid \theta') \cdot \mu_0(\theta') \cdot \mu_m(\theta) \cdot u_s(\hat{a}(\mu), \theta) \\ &= \sum_{\mu \in \text{support}(\tau)} \sum_{\theta \in \Theta} \sum_{m: \mu_m = \mu} \mu_0(\theta) \cdot q(m \mid \theta) \cdot u_s(\hat{a}(\mu), \theta) \\ &= \sum_{m \in M} \sum_{\theta \in \Theta} \mu_0(\theta) \cdot q(m \mid \theta) \cdot u_s(\hat{a}(\mu_m), \theta) = U_s(M, q). \end{aligned}$$

(2) Assume that τ is a Bayesian plausible distribution of posteriors with the expected utility u_s^* .

(3) Then we have

$$u_s^* = \sum_{\mu \in \text{support}(\tau)} \tau(\mu) \left[\sum_{\theta \in \Theta} \mu(\theta) \cdot u_s(\hat{a}(\mu), \theta) \right].$$

(4) Since Θ is finite, Carathéodory's theorem² implies that there exists a Bayesian plausible τ^* with finite support such that

$$u_s^* = \sum_{\mu \in \text{support}(\tau^*)} \tau^*(\mu) \left[\sum_{\theta \in \Theta} \mu(\theta) \cdot u_s(\hat{a}(\mu), \theta) \right].$$

(5) Define M so that $\text{support}(\tau^*) = \{\mu_m\}_{m \in M}$ and let $q(m \mid \theta) = \tau^*(\mu_m) \frac{\mu_m(\theta)}{\mu_0(\theta)}$.

(6) Clearly, $U_s(M, q) = U_s(\tau^*)$.

□

²In convex geometry Carathéodory's theorem states that if a point x of \mathbb{R}^d lies in the convex hull of a set P , there is a subset P' of P consisting of $d + 1$ or fewer points such that x lies in the convex hull of P' .

28. The key implication is that to evaluate whether the sender benefits from persuasion and to determine the value of an optimal information structure we need only ask how $U_s(\tau)$ varies over the space of Bayesian plausible distributions of posteriors.

29. Simplified problem:

$$\begin{aligned} & \underset{\tau}{\text{maximize}} && U_s(\tau) = \mathbf{E}_{\tau} U_s(\mu) \\ & \text{subject to} && \mathbf{E}_{\tau} \mu = \mu_0. \end{aligned}$$

5 Optimal information structure

30. The focus on sender-preferred equilibria implies that $U_s(\mu)$ is upper semicontinuous which in turn ensures the existence of an optimal information structure.

31. Let \hat{U}_s be the concave closure of the function $U_s(\mu): \Delta(\Theta) \rightarrow \mathbb{R}$:

$$\hat{U}_s(\mu) = \sup \{v \mid (\mu, v) \in \text{co}(U_s)\},$$

where $\text{co}(U_s)$ denotes the convex hull of the graph of U_s .

32. Clearly, $\hat{U}_s: \Delta(\Theta) \rightarrow \mathbb{R}$ is a concave function by construction. (check by yourself)

Actually, it is the smallest concave function that is everywhere weakly greater than U_s .

The following figure shows an example of the construction of \hat{U}_s . In the figure, the state space is binary, and we identify a distribution μ with the probability of one of the states.




Figure 2

33. Proposition: The expected utility of an optimal information structure is $\hat{U}_s(\mu_0)$, and revealing information structure is better than non-revealing if and only if $\hat{U}_s(\mu_0) > U_s(\mu_0)$.

Proof. (1) For any $(\mu', v) \in \text{co}(U_s)$, there exists a distribution of posteriors τ such that

$$(\mu', v) = \mathbf{E}_{\tau} (\mu, U_s(\mu)).$$

(2) Thus, $\text{co}(U_s)$ is the set of (μ', v) such that if the prior is μ' , there exists an information structure with the expected utility v .

(3) Hence, given the prior μ_0 , $\hat{U}_s(\mu_0)$ is the largest expected utility the sender can achieve with any information structure.

□

34. The following figure shows the function $U_s(\mu)$, the concave closure \hat{U}_s , and the optimal information structure for Example 10. In the figure, μ denotes the probability that the state is *guilty*.




Figure 3

- (a) Since there are two states *guilty* and *innocent*, each posterior belief can be represented by a number $\mu \in [0, 1]$, which denotes the probability being *guilty*.

- (b) The optimal action of receiver is

$$\hat{a}(\mu) = \begin{cases} \text{convict}, & \text{if } \mu \geq \frac{1}{2}, \\ \text{acquit}, & \text{if } \mu < \frac{1}{2}. \end{cases}$$

- (c) Thus, the prosecutor's expected utility U_s is a step function:

$$U_s(\mu) = \begin{cases} 1, & \text{if } \mu \geq \frac{1}{2} \text{ (the judge will choose } \textit{acquit}), \\ 0, & \text{if } \mu < \frac{1}{2} \text{ (the judge will choose } \textit{convict}). \end{cases}$$

- (d) The concave closure \hat{U}_s is

$$\hat{U}_s(\mu) = \begin{cases} 2\mu, & \text{if } \mu \leq \frac{1}{2}, \\ 1, & \text{if } \mu > \frac{1}{2}. \end{cases}$$

- (e) It is clear $\hat{U}_s(\mu_0) > U_s(\mu_0)$ and the expected utility of the optimal information structure is $\hat{U}_s(\mu_0)$. The prosecutor can benefit from persuasion if and only if $\mu_0 < \frac{1}{2}$.

- (f) By simple calculation, the optimal information structure induces the distribution of posteriors τ^* :

$$\tau^* = \frac{2}{5}\delta_0 + \frac{3}{5}\delta_{\frac{1}{2}}.$$

- 基于 $\mu_0 = 0.3$, 构建一个 distribution of posteriors, 并保持无偏条件成立。最简单构造方法：从 μ_0 向两边拉开, 分别得到后验猜测 μ_1 (概率 $\tau(\mu_1)$) 和 μ_2 (概率 $\tau(\mu_2)$), 保持加权平均值等于 μ_0 。比如, 后验猜测 $\mu_1 = 0.2$ 的概率 $\frac{1}{2}$, 后验猜测 $\mu_2 = 0.4$ 的概率 $\frac{1}{2}$; 再比如, 后验猜测 $\mu_1 = 0.1$ 的概率 $\frac{1}{3}$, 后验猜测 $\mu_2 = 0.4$ 的概率 $\frac{2}{3}$ 。
- 稍微拉开之后, μ_1 的期望效用 $U(\mu_1)$ 是零, μ_2 的期望效用 $U(\mu_2)$ 也是零, 所以期望效用 $U(\tau) = \tau(\mu_1)U(\mu_1) + \tau(\mu_2)U(\mu_2)$ 也是零。
- 继续拉开, 使得 $\mu_2 = \frac{1}{2}$, 这时期望效用 $U(\mu_2) = 1$, 并且 $U(\tau) = \tau(\mu_1)U(\mu_1) + \tau(\mu_2)U(\mu_2)$ 将大于 0。
- 使得期望效用 $U(\tau)$ 最大的选择将是：拉开到 $\mu_1 = 0$ (对应于信号 “不推荐”) 和 $\mu_2 = \frac{1}{2}$ (对应于信号 “推荐”)。为了使无偏条件成立, 后验猜测 $\mu_1 = 0$ 出现的概率是 $\frac{2}{5}$, 后验猜测 $\mu_2 = \frac{1}{2}$

出现的概率是 $\frac{3}{5}$ 。

- 为什么向右拉到 $\mu_2 = \frac{1}{2}$ ？为什么不继续向右拉开？超过 $\frac{1}{2}$ 之后， μ_2 越大，基于无偏条件， $\tau(\mu_2)$ 越小（对应于“推荐”的无条件概率减小）， $U(\tau)$ 也将越小。
- 为什么向左拉开到 $\mu_1 = 0$ ？ μ_1 越小，基于无偏条件， $\tau(\mu_2)$ 将越大（对应于“推荐”的无条件概率增大）， $U(\tau)$ 也将越大。
- 说服的要旨，就要用最低的代价，建立自己的可信度，然后把可信度用到刀刃上。在机制中，首先保证“推荐”的后验猜测恰好使得领导“批准”，其次通过调整“不推荐”的后验猜测增加“推荐”的无条件概率。
- 这里体现一个关键的权衡：通过“我不推荐，项目一定不会成功”这个保证，预先地主动排除了一些项目，这样建立了推荐的可信度——因为不推荐是不利于项目经理的利益的，所以领导很容易可以相信项目经理的保证。然后，项目经理就可以利用“不推荐”所建立的可信度，反过来增加自己推荐的可信度。

(g) Let $M = \{i, g\}$ such that $\mu_i = \delta_0$ and $\mu_g = \delta_{\frac{1}{2}}$. Then

$$q(g | guilty) = \tau^*(\mu_g) \frac{\mu_g(guilty)}{\mu_0(guilty)} = \frac{3}{5} \frac{1/2}{0.3} = 1,$$

$$q(g | innocent) = \tau^*(\mu_g) \frac{\mu_g(innocent)}{\mu_0(innocent)} = \frac{3}{5} \frac{1/2}{0.7} = \frac{3}{7}.$$

$$0.5 = \text{Prob}(\text{成功} | \text{推荐}) = \frac{\text{Prob}(\text{推荐} | \text{成功}) \times \text{Prob}(\text{成功})}{\text{Prob}(\text{推荐})} = \frac{\text{Prob}(\text{推荐} | \text{成功}) \times 0.3}{0.6}$$

因此 $\text{Prob}(\text{推荐} | \text{成功}) = 1$ 。

$$0.5 = \text{Prob}(\text{失败} | \text{推荐}) = \frac{\text{Prob}(\text{推荐} | \text{失败}) \times \text{Prob}(\text{失败})}{\text{Prob}(\text{推荐})} = \frac{\text{Prob}(\text{推荐} | \text{失败}) \times 0.7}{0.6}$$

因此 $\text{Prob}(\text{推荐} | \text{失败}) = \frac{3}{7}$ 。

35. If $\mu_0 \geq 0.5$, then revealing information structure cannot be better than non-revealing.

36. Corollary: If $U_s(\mu)$ is concave, the sender does not benefit from persuasion for any prior. If $U_s(\mu)$ is convex and not concave, the sender benefits from persuasion for every prior.

Proof. The sender benefits from persuasion is and only if there exists a τ such that $\mathbf{E}_\tau[U_s(\mu)] > U_s(\mathbf{E}_\tau[\mu])$. \square

37. We say “there is information the sender would like to share” if there is a posterior that is better for the sender than the prior, that is, there exists μ such that

$$U_s(\mu) = \sum_{\theta \in \Theta} \mu(\theta) \cdot u_s(\hat{a}(\mu), \theta) > \sum_{\theta \in \Theta} \mu(\theta) \cdot u_s(\hat{a}(\mu_0), \theta).$$

In other words, there is a μ such that, if the sender had private information that led him to believe μ , he would prefer to share this information with the receiver rather than have the receiver act based on μ_0 .

We say the reserver’s preference is discrete at belief μ if the receiver’s expected utility from her preferred action $\hat{a}(\mu)$ is bounded away from her expected utility from any other action, i.e., if there is an $\epsilon > 0$ such that for any $a \neq \hat{a}(\mu)$, $\mathbf{E}_\mu u_r(\hat{a}(\mu), \theta) > \mathbf{E}_\mu u_r(a, \theta) + \epsilon$.

38. Proposition: If there is no information the sender would share, the sender does not benefit from persuasion. If there is information the sender would share and the receiver's preference is discrete at the prior, the sender benefits from persuasion.

Proof. (1) If there is no information the sender would share, then for any information structure which induces a distribution of posteriors τ ,

$$\mathbf{E}_\tau U_s(\mu) \leq \underbrace{\mathbf{E}_\tau \mathbf{E}_\mu u_s(\hat{a}(\mu_0), \theta)}_{\text{Bayesian plausibility}} = U_s(\mu_0).$$

Informally, any realization of message m leads the receiver to take an action $\hat{a}(\mu_m)$ the sender weakly dislikes relative to the default action $\hat{a}(\mu_0)$. Hence, a completely noninformative information structure is optimal.

- (2) Since the receiver's preference is discrete at the prior μ_0 , there exists an $\epsilon > 0$ such that for any $a \neq \hat{a}(\mu_0)$, $\sum_\theta \mu_0(\theta) \cdot u_r(\hat{a}(\mu_0), \theta) > \sum_\theta \mu_0(\theta) \cdot u_r(a, \theta) + \epsilon$.
- (3) Since $u_r(a, \theta)$ is continuous in θ , $\sum_\theta \mu(\theta) \cdot u_r(a, \theta)$ is continuous in μ .
- (4) Thus, there is a $\delta > 0$ such that for any $\mu \in B_\delta(\mu_0)$ and for any $a \neq \hat{a}(\mu_0)$, $\sum_\theta \mu(\theta) \cdot u_r(\hat{a}(\mu_0), \theta) > \sum_\theta \mu(\theta) \cdot u_r(a, \theta)$.
- (5) Hence, $\hat{a}(\mu) = \hat{a}(\mu_0)$ for any $\mu \in B_\delta(\mu_0)$.
- (6) Since there is information the sender would share, there exists μ_h such that $\sum_\theta \mu_h(\theta) \cdot u_s(\hat{a}(\mu_h), \theta) > \sum_\theta \mu_h(\theta) \cdot u_s(\hat{a}(\mu_0), \theta)$.
- (7) Consider a ray from μ_h through μ_0 . Since μ_0 is not on the boundary of $\Delta(\Theta)$, there exists a belief μ_ℓ on that ray such that $\mu_\ell \in B_\delta(\mu_0)$ and $\mu_0 = \gamma\mu_\ell + (1 - \gamma)\mu_h$ for some $\gamma \in (0, 1)$.
- (8) Consider the Bayesian plausible distribution of posteriors $\tau = \gamma\mu_\ell + (1 - \gamma)\mu_h$.
- (9) We have

$$\begin{aligned} \mathbf{E}_\tau \mathbf{E}_\mu u_s(\hat{a}(\mu), \theta) &= \gamma U_s(\mu_\ell) + (1 - \gamma) U_s(\mu_h) \\ &> \gamma \sum_\theta \mu_\ell(\theta) \cdot u_s(\hat{a}(\mu_0), \theta) + (1 - \gamma) \sum_\theta \mu_h(\theta) \cdot u_s(\hat{a}(\mu_0), \theta) = U_s(\mu_0). \end{aligned}$$

Therefore, the sender benefits from persuasion. □

39. Lemma: If A is finite, the receiver's preference is discrete at the prior generically. □

Proof. Omitted. □

40. Application: Lobbying.

Consider a setting where a lobbying group commissions a study with the goal of influencing a benevolent, but nonetheless rational, politician. The politician (Receiver) chooses a unidimensional policy $a \in [0, 1]$. The state $\theta \in [0, 1]$ is the socially optimal policy. The lobbyist (Sender) is employed by the interest group whose preferred action is $a_0 = \alpha\theta + (1 - \alpha)\theta_0$ with $\alpha \in [0, 1]$ and $\theta_0 > 1$. Politician's payoff $-(a - \theta)^2$ and lobbyist's payoff $-(a - a_0)^2$.

Since politician's payoff is $-(a - \theta)^2$, $\hat{a}(\mu) = \mathbf{E}_\mu[\theta]$. Given this \hat{a} , we have

$$U_s(\mu) = -(1 - \alpha)^2 \theta_0^2 + 2(1 - \alpha)^2 \theta_0 \mathbf{E}_\mu[\theta] - \alpha^2 \mathbf{E}_\mu[\theta^2] + (2\alpha - 1)(\mathbf{E}_\mu[\theta])^2.$$

U_s is linear in μ when $\alpha = \frac{1}{2}$, strictly convex when $\alpha > \frac{1}{2}$, and strictly concave when $\alpha < \frac{1}{2}$.

Therefore we have full disclosure if $\alpha > \frac{1}{2}$ and no disclosure if $\alpha < \frac{1}{2}$. There is thus a natural sense in which some alignment of preferences is necessary for information to be communicated in equilibrium even when Sender has the ability to commit.

Note that the lobbyist either commissions a fully revealing study or no study at all.

The optimal information structure is independent of θ_0 . This is important because θ_0 also captures a form of disagreement between the lobbyist and the politician. We might have expected communication to be difficult when θ_0 is much greater than one. Unlike α , however, θ_0 does not affect the way the lobbyist's payoff varies across realizations of a message. The loss the lobbyist suffers from high values of θ_0 is thus a sunk cost and does not affect the decision of how best to persuade.

Task

- Reading:
- Understanding: