

Game Theory Principles V

Andreas Blume
Department of Economics
University of Pittsburgh
Pittsburgh, PA 15260

Lessons from Experiment II

1. In a class of games some players may have private information.
2. Other players may have probability assessments of this private information.
3. Players may update their probability assessments over time, as new information arrives.

4. We need a general way of accounting for private information in games.
5. We need to understand how rational agents update probability assessments.

A second example of private information in games:

Policy advice by an informed expert

1. Consider a decision maker who must choose between three policies a , b and c .
2. The decision maker is uncertain about the payoffs that will result from the three policies.
3. The decision maker consults an informed expert.
4. The expert knows the payoff consequences of all actions.

5. We will refer to the expert's private information as her **type**
6. Suppose the expert has two possible types, s and t .
7. Suppose that both types of the expert are equally likely.
8. Suppose that that the decision maker and the expert have identical preferences.

9. Suppose that a is the payoff-maximizing action if the expert's private information, i.e. her type, is s .
10. Suppose that b is the payoff-maximizing action if the expert's type is t .
11. Suppose that if the decision maker believes that both types are equally likely, his payoff maximizing action is c .

The following payoff matrix has payoff entries that conform with our assumptions.

	<i>a</i>	<i>b</i>	<i>c</i>
<i>s</i>	3, 3	0, 0	2, 2
<i>t</i>	0, 0	3, 3	2, 2

Each cell of the payoff matrix contains first the expert's and second the policy maker's payoff from the corresponding action given the corresponding type; e.g., the expert's payoff for type *t* equals 0 when action *a* is taken and it equals 3 when action *b* is taken.

Recall that both types are equally likely: $\Pr(s) = \Pr(t) = \frac{1}{2}$.

The above payoff matrix does not (!) represent a game.

Policy advice game

1. The expert privately learns her type, s or t .
2. After learning her type, the expert sends a message, m or n , to the policy maker.
3. After observing the message, the policy maker takes an action, a , b or c .
4. Finally, payoffs are realized.

Strategies in the policy advice game

1. Recall that in any game a strategy is a **complete contingent plan**.
2. The expert's strategies are all functions, σ , from her possible types into messages.
3. The policy maker's strategies are all functions, ρ , from messages into actions.
4. You should verify that the expert has four strategies and the policy maker has nine strategies.

An unappealing Nash equilibrium of the policy advice game

1. Expert's strategy: $\sigma(s) = \sigma(t) = n$
2. Policy maker's strategy: $\rho(m) = \rho(n) = c$
3. In this equilibrium the expert ignores her information and always sends the same message.
4. Likewise, the policy maker ignores the expert's message and always takes the same action.

5. Note that if the policy maker ignores the expert's message, the expert cannot gain by deviating to a different message, for any of his types.
6. Note that if both types of the expert send the same message, the message is uninformative about the expert's type. Therefore both types remain equally likely. As long as both types are equally likely, action c maximizes the policy maker's payoff.

A more appealing Nash equilibrium of the policy advice game

1. Expert's strategy: $\sigma(s) = m$ and $\sigma(t) = n$
2. Policy maker's strategy: $\rho(m) = a$ and $\rho(n) = b$
3. If only type s sends message m , the message is informative and indicates to the policy maker that the true type is s .
4. Upon upon receiving message m , the policy maker **updates his belief** from $\Pr(s) = \frac{1}{2}$ to $\Pr(s) = 1$.
5. Given the updated belief $\Pr(s) = 1$, the unique best reply of the receiver is to take the action a .

In the appealing (separating) equilibrium of the policy advice game, belief updating is straightforward.

In the Monty Hall game from Experiment II, correct belief updating is somewhat less transparent.

The general procedure for proper belief updating that is always valid is known as **Bayes's Rule**

In order to derive Bayes's rule, we need to introduce a few basic ideas about probabilities.

Probability Facts

Throughout we will refer to the following example:

1. Pairs of tosses of fair coins.
2. A coin is fair if heads, H , and tails, T , are equally likely.
3. The pair of coin tosses results in one of the following four outcomes: HH , HT , TH and TT .

We describe a random process using:

1. a set of basic outcome, called the **sample space** S ,
(in the example: $S = \{HH, HT, TH, TT\}$)
2. subsets of the sample space are called **events**,
(in the example one event A is for at least one coin flip in a pair to come up heads, $A = \{HH, HT, TH\}$, another event B is for at least one of the two coins to come up tails, $B = \{HT, TH, TT\}$.)

3. The joint event “ A and B ” is simply the intersection $A \cap B$.
4. The event “ A or B ” is the union $A \cup B$.
5. The event “not B ” is denoted \overline{B} .

The probability of an event A , denoted $\Pr(A)$ is a number that satisfies $0 \leq \Pr(A) \leq 1$.

If A and B are mutually disjoint events, i.e. $A \cap B = \emptyset$, then $\Pr(A \cup B) = \Pr(A) + \Pr(B)$. (in the example the events “both coins come up heads” and “both coins come up tails” are mutually exclusive.)

The events A and B are **independent** if the probability of the joint event $A \cap B$, i.e. both A and B occur, is the product of the probabilities of event A and event B , i.e. $\Pr(A \cap B) = \Pr(A) \cdot \Pr(B)$.

In the example the events “the first coin comes up heads” and “the second coin comes up tails” are independent. The events “at least one H ” and “at least one T ” are not independent.

Conditional Probability

The conditional probability of event A occurring conditional on event B having occurred is defined as

$$\Pr(A|B) = \frac{\Pr(A \cap B)}{\Pr(B)}.$$

In the example, the probability of “at least one H ” conditional on “at least one T ” equals:

$$\begin{aligned} & \Pr(\text{“at least one } H\text{”} \mid \text{“at least one } T\text{”}) \\ &= \frac{\Pr(\text{“at least one } H\text{”} \cap \text{“at least one } T\text{”})}{\Pr(\text{“at least one } T\text{”})} \\ &= \frac{\Pr(\{HT, TH\})}{\Pr(\{HT, TH, TT\})} = \frac{1/2}{3/4} = \frac{2}{3}. \end{aligned}$$

Bayes's Rule (also known as Bayes' Rule, but it's time to change ...)

The definition for conditional probability can be rewritten as follows:

$$\Pr(A \cap B) = \Pr(A|B)\Pr(B).$$

Exchanging the roles of A and B , we also have

$$\Pr(A \cap B) = \Pr(B|A)\Pr(A).$$

Combine both expressions:

$$\Pr(A|B)\Pr(B) = \Pr(B|A)\Pr(A),$$

and divide by $\Pr(B)$ to obtain ...

BAYES'S RULE

$$\Pr(A|B) = \frac{\Pr(B|A)\Pr(A)}{\Pr(B)}$$

Preparing to rewrite Bayes's rule:

Since $B \cap A$ and $B \cap \bar{A}$ are disjoint, we have

$$\begin{aligned}\Pr(B) &= \Pr((B \cap A) \cup (B \cap \bar{A})) \\ &= \Pr(B \cap A) + \Pr(B \cap \bar{A}) \\ &= \Pr(B|A)\Pr(A) + \Pr(B|\bar{A})\Pr(\bar{A})\end{aligned}$$

Hence, we get the following version of ...

BAYES'S RULE

$$\Pr(A|B) = \frac{\Pr(B|A)\Pr(A)}{\Pr(B|A)\Pr(A) + \Pr(B|\bar{A})\Pr(\bar{A})}$$

A trivial application of Bayes's rule

Recall the appealing equilibrium in the policy advice game:

1. Type s sends message m and type t sends message n .
2. Initially types s and t are equally likely, i.e. $\Pr(s) = \Pr(t) = \frac{1}{2}$. These are the policy maker's **prior beliefs**.
3. We want to know the policy maker's **posterior belief** $\Pr(s|m)$ conditional on having received message m .

4. According to Bayes's rule:

$$\begin{aligned}\Pr(s|m) &= \frac{\Pr(m|s)\Pr(s)}{\Pr(m)} \\ &= \frac{\Pr(m|s)\Pr(s)}{\Pr(m|s)\Pr(s) + \Pr(m|t)\Pr(t)} \\ &= \frac{1 \cdot \frac{1}{2}}{1 \cdot \frac{1}{2} + 0 \cdot \frac{1}{2}} \\ &= 1\end{aligned}$$

Thus Bayes's rule confirms what we already knew from simple intuition: Given that only type s sends message m , if the policy maker observes message m he knows that the expert's type is s .

Things are not always so simple. Consider the modified policy advice game:

	<i>a</i>	<i>b</i>	<i>c</i>
<i>s</i>	3, 3	0, 0	2, 2
<i>t</i>	3, 0	3, 3	2, 2

Now type *t* of the expert is just as happy with action *a* as with action *b*.

1. Suppose the policy maker continues responding with action a to message m and with action b to message n .
2. Also, suppose the expert continues sending message m as type s .
3. Unlike before, suppose now that expert of type t randomizes over both messages m and n . For the sake of concreteness suppose she sends both messages with equal probability.

We can write the strategy of the expert follows:

1. The expert's probability of sending message m conditional on being type t equals: $\sigma(m|t) = \frac{1}{2}$.
2. The expert's probability of sending message n conditional on being type t equals: $\sigma(n|t) = \frac{1}{2}$.
3. The expert's probability of sending message m conditional on being type s equals: $\sigma(m|s) = 1$.
4. The expert's probability of sending message n conditional on being type s equals: $\sigma(n|s) = 0$.

What is the posterior probability of type s conditional on observing message m ?

Observe that now there is some chance that type t sends message m . Hence the posterior probability of type s given message m , denoted $\mu(s|m)$ is now less than one.

What exactly is $\mu(s|m)$?

Bayes's Rule to the rescue:

$$\begin{aligned}\mu(s|m) &= \frac{\sigma(m|s) \cdot \Pr(s)}{\sigma(m|s) \cdot \Pr(s) + \sigma(m|t) \cdot \Pr(t)} \\ &= \frac{1 \cdot \frac{1}{2}}{1 \cdot \frac{1}{2} + \frac{1}{2} \cdot \frac{1}{2}} \\ &= \frac{2}{3}\end{aligned}$$

Therefore conditional on receiving message m , action a remains a best response for the receiver. Hence, the postulated strategies form a Nash equilibrium.

Observe that

1. in this equilibrium the receiver's payoff is lower than before,
2. if the type t sender sent the message m with any higher probability, the equilibrium would collapse,
3. if we further increased the payoff of the type t sender from action a , there would be no informative equilibrium, i.e. **a substantial bias in the expert's preferences undermines information transmission.**

The Monty Hall Game

(Analysis of Experiment II)

Define the following events:

1. L := “the token is in L”
2. R := “the token is in R”
3. M := “the token is in M”
4. “ L ” := “Monty reveals that L is empty”
5. “ R ” := “Monty reveals that R is empty”
6. “ M ” := “Monty reveals that M is empty”

Without loss of generality, suppose that the initial choice is R .

With slight loss of generality, but without affecting the final conclusion, suppose that Monty chooses with equal probability between M and L in the event “the token is in R .”

We are interested in the probability that the token is in L given that Monty revealed that M is empty, denoted

$$\Pr(L | \text{“}M\text{”}).$$

If this probability is greater than one half, it is preferable to switch from R (the original choice) to L (the other available choice).

According to Bayes's Rule:

$$\begin{aligned}\Pr(L|\text{"M"}) &= \frac{\Pr(\text{"M"}|L)\Pr(L)}{\Pr(\text{"M"}|L)\Pr(L) + \Pr(\text{"M"}|R)\Pr(R)} \\ &= \frac{1 \cdot \frac{1}{3}}{1 \cdot \frac{1}{3} + \frac{1}{2} \cdot \frac{1}{3}} \\ &= \frac{2}{3}\end{aligned}$$

That is, switching is indeed optimal.