“Game Theory"

Yves Zenou
Stockholm University
Office: A781
Email: yves.zenou@ne.su.se
Office hours: Tuesdays 16.00-17.00

Lecture Note 6

- Bayesian games
- Information
- Signals and types
- Bayesian Nash equilibrium
- Bayes' rule

Why bother?

- In reality, do agents have perfect information about the game in which they participate?
  - Nations/political groupings may be uncertain over the military strength of their opponents and/or their preferences over outcomes.
  - Firms may not know their competitors' progress in R&D.
  - Political parties may not know if their opponents (and or the media) know about their scandals...
- Hence, reality is full of uncertainty – knowing how to model these kinds of situations may be helpful...

What is a Bayesian game?

- Generalizes the strategic game by allowing players to be imperfectly informed about some aspect in his/her environment that may influence his/her choice of action.
- Note! Preferences are assumed to be represented by vNM payoff functions (the players need to be able to rank different ”lotteries”).

Battle of the sexes: Multiple NE (Coordination games)

Partner 2

Partner 2		*B*	*S*
	B	2, 1	0, 0
	S	0, 0	1, 2

Battle of sexes with imperfect information (Osborne)

- Player 1 is unsure whether player 2 prefers to go out with her or prefers to avoid her, whereas player 2, as before, knows player 1’s preferences.
- In other words, Player 1 is unsure of what type player 2 is.
- Player 2 knows what type he/she is and also player 1’s type.

Battle of sexes with imperfect information

- Specifically, suppose player 1 thinks that with probability ½ player 2 wants to got out with her, and with probability ½ player 2 wants to avoid her.

Battle of sexes with imperfect information

- That is, player 1 thinks that with probability ½ she is playing the game on the left of this Figure and with probability ½ she is playing the game on the right.

1

Prob 1/2

2

2 wishes to meet with 1

	B	*S*
B	2, 1	0, 0
S	0, 0	1, 2

Prob 1/2

2

2 wishes to avoid 1

	B	*S*
B	2, 0	0, 2
S	0, 1	1, 0

Battle of sexes with imperfect information

- We can think of there being two *states*, one in which the players' payoffs are given in the left table and one in which these payoffs are given in the right table.
- Player 2 knows the state—she knows whether she wishes to meet or avoid player 2, whereas player 1 does not; player 1 assigns probability ½ to each state.

Nash equilibrium

- The notion of Nash equilibrium for a strategic game models a state in which each player's beliefs about the other players' actions are correct, and each player acts optimally, given her beliefs.
- We wish to generalize this notion to the current situation.

Battle of sexes with imperfect information

- From player 1's point of view, player 2 has two possible *types*, one whose preferences are given in the left table, and one whose preferences are given in the right table.
- Player 1 does not know player 2's type, so to choose an action rationally she needs to form a **belief** about the action of each type.
- Given these beliefs and her belief about the likelihood of each type, she can calculate her expected payoff to each of her actions.

Battle of sexes with imperfect information

- For example, if player 1 thinks that the type who wishes to meet her will choose *B* and the type who wishes to avoid her will choose *S*,
- then she thinks that playing *B* will yield her a payoff of 2 with probability ½ and a payoff of 0 with probability ½, so that her **expected payoff** of playing B is: ½ x 2 + ½ x 0 = 1
- Similarly, her expected payoff of playing *S* is:

  ½ x 0 + ½ x 1 = ½

- Similar calculations for the other combinations of actions for the two types of player 2 yield the following expected payoffs:

	(B,B)	(B,S)	(S,B)	(S,S)
B	2	1	1	0
S	0	1/2	1/2	1

Battle of sexes with imperfect information

- For this situation, we define a pure strategy *Nash equilibrium* to be a triple of actions, one for player 1 and one for each type of player 2, with the property that:
- (*i*) the action of player 1 is optimal, given the actions of the two types of player 2 (and player 1's belief about the state)
- (*ii*) the action of each type of player 2 is optimal, given the action of player 1.

Battle of sexes with imperfect information

- That is, we treat the two types of player 2 as separate players, and analyze the situation as a three-player strategic game.
- What is the unique NE of this game?

- Player 2 is the type who whishes to meet player 1.
- Player 3 is the type who whishes to avoid player 1.

Battle of sexes with imperfect information

- Normal form: first payoff is for player 1, the second for player 2 of type y2 and third for n2.

	(B,B)	(B,S)	(S,B)	(S,S)
B	(2,1,0)	(1,1,2)	(1,0,0)	(0,0,2)
S	(0,0,1)	(½,0,0)	(½,2,0)	(1,2,0)

Normal form

(P3) B			(P3) S		
	(P2) B	(P2) S		(P2) B	(P2) S
(P1)B	(2,1,0)	(1,0,0)	(P1)B	(1,1,2)	(0,0,2)
(P1)S	(0,0,1)	(½,2,0)	(P1)S	(½,0,0)	(1,2,0)

BR functions

	(P3) B		(P3) S	
	(P2) B	(P2) S	(P2) B	(P2) S
(P1)B	(**2**,**1**,0)	(**1**,0,0)	(**1**,**1**,**2**)	(0,0,**2**)
(P1)S	(0,0,**1**)	(½,**2**,**0**)	(½,0,0)	(**1**,**2**,**0**)

Two NE

(P3) B			(P3) S		
	(P2) B	(P2) S		(P2) B	(P2) S
(P1)B	(**2**,**1**,0)	(**1**,0,0)	(P1)B	(**1**,**1**,**2**)	(0,0,**2**)
(P1)S	(0,0,**1**)	(½,**2**,**0**)	(P1)S	(½,0,0)	(**1**,**2**,**0**)

- We can interpret the actions of the two types of player 2 to reflect player 2's intentions in the hypothetical situation *before* she knows the state.
- We can tell the following story.
- Initially player 2 does not know the state; she is informed of the state by a *signal* that depends on the state. Before receiving this signal, she plans an action for each possible signal.
- After receiving the signal she carries out her planned action for that signal.

- We can tell a similar story for player 1.
- To be consistent with her not knowing the state when she takes an action, her signal must be **uninformative**: it must be the same in each state. Given her signal, player 1 is unsure of the state; when choosing an action she takes into account her belief about the likelihood of each state, given her signal.

Variant of the Battle of Sexes

- Consider another variant of the situation modeled by the *BoS*, in which neither player knows whether the other wants to go out with her.
- Specifically, suppose that player 1 thinks that with probability ½ , player 2 wants to go out with her, and with probability ½ , player 2 wants to avoid her
- Player 2 thinks that with probability 2/3, player 1 wants to go out with her and with probability 1/3, player 1 wants to avoid her.
- As before, assume that each player knows her own preferences.

Battle of sexes with imperfect information

- We can model this situation by introducing four states, one for each of the possible configurations of preferences.
- We refer to these states as *yy* (each player wants to go out with the other), *yn* (player 1 wants to go out with player 2, but player 2 wants to avoid player 1), *ny*, and *nn*.
- Note: *y* stands for ”yes” and *n* for ”no”.

Example (ctd.)

Type y_1 of player 1 believes that the probability of each of the states yy and yn is $\frac{1}{2}$; type n_1 of player 1 believes that the probability of each of the states ny and nn is $\frac{1}{2}$. Similarly, type y_2 of player 2 believes that the probability of state yy is $\frac{2}{3}$ and that of state ny is $\frac{1}{3}$; type n_2 of player 2 believes that the probability of state yn is $\frac{2}{3}$ and that of state nn is $\frac{1}{3}$. This model of the situation is illustrated in Figure 275.1.

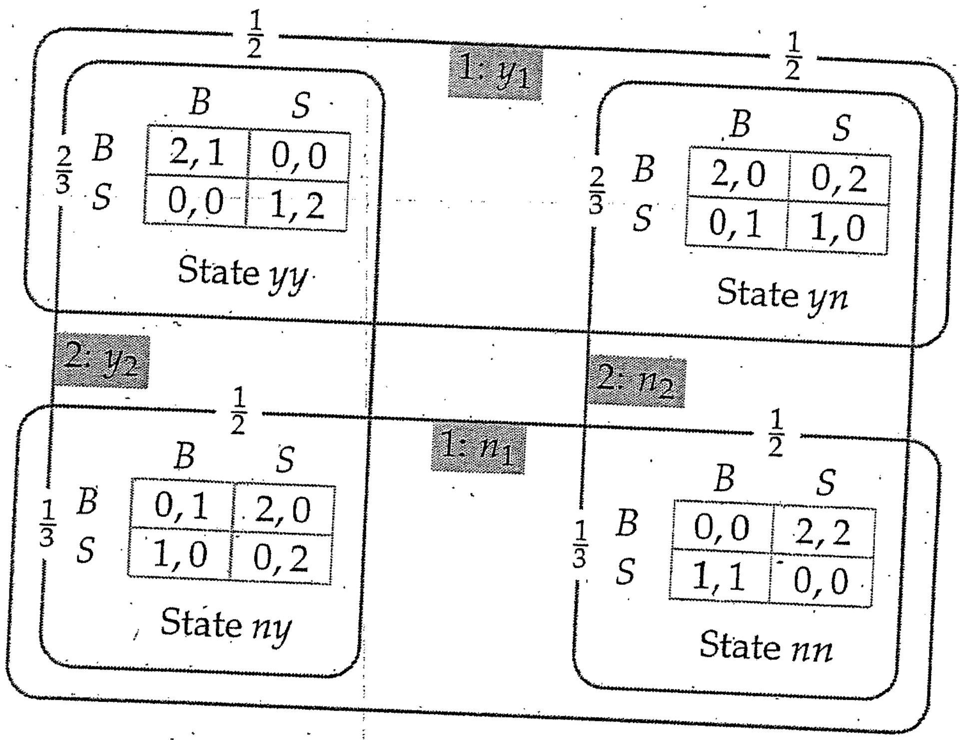

Figure 275.1 A variant of *BoS* in which each player is unsure of the other player's preferences. The frame labeled i: x encloses the states that generate the signal x for player i; the numbers printed over this frame next to each table are the probabilities that type x of player i assigns to each state that she regards to be possible.

Battle of sexes with imperfect information

- As in the previous example, to study the equilibria of this model we consider the players' plans of action before they receive their signals.
- That is, each player plans an action for each of the two possible signals she may receive.
- We may think of there being four players: the two types of player 1 and the two types of player 2.

Battle of sexes with imperfect information

- A *Nash equilibrium* consists of four actions, one for each of these players, such that the action of each type of each original player is optimal, given her belief about the state after observing her signal, and given the actions of each type of the other original player.

Battle of sexes with imperfect information

- Caculate the expected payoffs of each player.

- Expected payoff of **type y1 of player 1**:
- Each row corresponds to a pair of actions for the two types of player 2; the action of type y2 is listed first, that of type n2, second.

	(B,B)	(B,S)	(S,B)	(S,S)
B	2	1	1	0
S	0	1/2	1/2	1

- Expected payoff of **type n1 of player 1**:

	(B,B)	(B,S)	(S,B)	(S,S)
B	0	1	1	2
S	1	1/2	1/2	0

- Expected payoff of **type y2 of player 2**:
- Each row corresponds to a pair of actions for the two types of player 1; the action of type y1 is listed first, that of type n1, second.

	(B,B)	(B,S)	(S,B)	(S,S)
B	1	2/3	1/3	0
S	0	2/3	4/3	2

- Expected payoff of **type n2 of player 2**:

	(B,B)	(B,S)	(S,B)	(S,S)
B	0	1/3	2/3	1
S	2	4/3	2/3	0

Battle of sexes with imperfect information

- Determine the NE of this game.

	(y1,n1) (B,B)	(y1,n1) (B,S)	(y1,n1) (S,B)	(y1,n1) (S,S)
y2,n2 (B,B)	y1,n1,y2,n2 2, 0, 1, 0	2,1,2/3,1/3	0,0,1/3,2/3	0,1,0,1
(B,S)	1,1,1,2	1,1,2/3,4/3	1/2,1,1/3,4/3	½, ½,0,0
(S,B)	1,1,0,0	1,1,2/3,1/3	1/2,1,4/3,2/3	½, ½,2,1
(S,S)	0,2,0,2	0,0,2/3,4/3	1,2,4/3,2/3	1,0,2,0

	(y1,n1) (B,B)	(y1,n1) (B,S)	(y1,n1) (S,B)	(y1,n1) (S,S)
y2,n2 (B,B)	y1,n1,y2,n2 **2**, 0, **1**, 0	**2**,**1**,**2/3**,1/3	0,0,1/3,2/3	0,**1**,0,**1**
(B,S)	**1**,**1**,**1**,**2**	**1**,**1**,**2/3**,**4/3**	1/2,**1**,1/3,**4/3**	½, ½,0,0
(S,B)	**1**,**1**,0,0	**1**,**1**,**2/3**,1/3	1/2,**1**,**4/3**,**2/3**	½, ½,**2**,**1**
(S,S)	0,**2**,0,**2**	0,0,**2/3**,**4/3**	**1**,**2**,**4/3**,**2/3**	**1**,0,**2**,0

	(y1,n1) (B,B)	(y1,n1) (B,S)	(y1,n1) (S,B)	(y1,n1) (S,S)
y2,n2 (B,B)	y1,n1,y2,n2 **2**, 0, **1**, 0	**2**,**1**,**2/3**,1/3	0,0,1/3,2/3	0,**1**,0,**1**
(B,S)	**1**,**1**,**1**,**2**	**1**,**1**,**2/3**,**4/3**	1/2,**1**,1/3,**4/3**	½, ½,0,0
(S,B)	**1**,**1**,0,0	**1**,**1**,**2/3**,1/3	1/2,**1**,**4/3**,**2/3**	½, ½,**2**,**1**
(S,S)	0,**2**,0,**2**	0,0,**2/3**,**4/3**	**1**,**2**,**4/3**,**2/3**	**1**,0,**2**,0

General definition of Bayesian games

- At the start of the game a state is realized. The players do not observe this state.
- Rather, each player receives a *signal* that may give her some information about the state.
- Denote the signal player *i* receives in state ω by: $\tau_i(\omega)$
- The function τ_i is called player *i*'s *signal function*.

General definition of Bayesian games

- The states that generate any given signal t_i are said to be *consistent* with t_i .

- The sizes of the sets of states consistent with each of player *i*'s signals reflect the quality of player *i*'s information.

- If, for example, $\tau_i(\omega)$ is different for each value of ω, then player *i* knows, given, her signal, the state that has occurred; after receiving her signal, she is perfectly informed about all the players' relevant characteristics.
- At the other extreme, if $\tau_i(\omega)$ is the same for all states, then player *i*'s signal conveys no information about the state.

- If $\tau_i(\omega)$ is constant over some subsets of the set of states, but is not the same for all states, then player *i*'s signal conveys partial information

- For example, if there are three states, *ω*1, *ω*2, and *ω*3, and $\tau_i(\omega_1) \neq \tau_i(\omega_2) = \tau_i(\omega_3)$, then when the state is *ω*1 player *i* knows that it is *ω*1, whereas when it is either *ω*2 or *ω*3, she knows only that it is one of these two states.

- We refer to player *i* in the event that she receives the signal t_i as **type** t_i of player *i*.
- Each type of each player holds a **belief** about the likelihood of the states consistent with her signal.
- If, for example, $t_i = \tau_i(\omega_1) = \tau_i(\omega_2)$, then type t_i of player *i* assigns probabilities to *ω*1 and *ω*2.

- A player who receives a signal consistent with only one state naturally assigns probability 1 to that state.

Definition: Bayesian game

- A Bayesian game consists of:
- A set of **players**, *N*={1,...,*n*)
- A set of **states**, *Ω*={*ω*1*, ω*2*,..., ω*S}

and for each player *i*

- A set of **actions** *A*={a1*,a2,..., a*n}
- A set of **signals,** $T = \{\tau(\omega_1),...,\tau(\omega_M)\}$
- where $\tau(\omega)$ is the signal function that associates a signal with each state
- For each signal that she may receive, a **belief** about the states consistent with the signal (a probability distribution over the set of states which which the signal is associated).
- A **Bernoulli payoff function** over pairs of *(a, ω),* that is, over all action profiles in all possible states.

Definition: Bayesian game

- Note that the set of actions of each player is independent of the state.
- Each player may care about the state, but the state of actions available to her is the same in every state.

Example of Bayesian Game: BoS

- Example when Player 1 is unsure of what type player 2 is but Player 2 knows what type he/she is and also player 1's type.
- **Players:** The pair of people.
- **States:** The set of states is *{meet*, *avoid}*.
- **Actions:** Set of actions of each player is *{B*, *S}*.
- **Signals:** Player 1 may receive a single signal, say *z*; her signal function τ_1 satisfies:

$$\tau_1(meet) = \tau_1(avoid) = z$$

- Player 2 receives one of two signals, say *m* and *v*; her signal function τ_2 satisfies $\tau_2(avoid) = v$ and $\tau_2(meet) = m$

- **Beliefs:** Player 1 assigns probability ½ to each state after receiving the signal *z*.
- Player 2 assigns probability 1 to the state *meet* after receiving the signal *m*, and probability 1 to the state *avoid* after receiving the signal *v*.

- **Payoffs:** The payoffs *ui*(*a*,*meet*) of each player *i* for all possible action pairs are given in the left panel of the Figure, and the payoffs *ui*(*a*, *avoid*) are given in the right panel.

Example of Bayesian Game: BoS

- Example when Player 1 is unsure of what type player 2 is and Player 2 is unsure of player 1's type.

- **Players:** The pair of people.
- **States:** The set of states is *{yy*, *yn*, *ny*, *nn}*.
- **Actions:** Set of actions of each player is *{B*, *S}*.

- **Signals:** Player 1 may receive one of two signals, *y*1 and *n*1; her signal function τ_1 satisfies:

$$\tau_1(yy) = \tau_1(yn) = y_1 \quad \text{and} \quad \tau_1(ny) = \tau_1(nn) = n_1$$

- Player 2 receives one of two signals, *y*2 and *n*2; her signal function τ_2 satisfies:

$$\tau_2(yy) = \tau_1(ny) = y_2 \quad \text{and} \quad \tau_2(yn) = \tau_2(nn) = n_2$$

- **Beliefs:** Player 1 assigns probability ½ to each of the states *yy* and *yn* after receiving the signal *y*1, and probability ½ to each of the states *ny* and *nn* after receiving the signal *n*1

- Player 2 assigns probability 2/3 to the state *yy* and probability 1/3 to the state *ny* after receiving the signal *y*2,

  and probability 2/3 to the state *yn* and probability 1/3 to the state *nn* after receiving the signal *n*2.

- **Payoffs:** The payoffs *ui*(*a*,*ω*) of each player *i* for all possible action pairs and states are given in the Figure.

- In a strategic game, each player chooses an action.
- In a Bayesian game, each player chooses a collection of actions, one for each signal she may receive.
- That is, in a Bayesian game **each type of each player** chooses an action.

Sequence of things...

- First, a state is realized for the game (the players do not know which state).
- Second, the players receive signals that give them some information on which state has been realized (this information may be ”zero”, partial or full where the last case means that the player knows exactly which state has been realized, i.e. full information.
- Third, the players choose their actions given their beliefs, maximizing expected utility according to;

$$\text{Expected payoff of type } t_i \text{ of player } i \text{ when choosing action } a_i \text{ is;}$$

$$\sum_{\omega \in \Omega} \Pr(\omega \text{ given } t_i) u_i((a_i, \hat{a}_{-i}(\omega)), \omega)$$

- **DEFINITION**: A **Nash equilibrium of a Bayesian game** is a Nash equilibrium of the strategic game (with vNM preferences) defined as follows:
- **Players:** The set of all pairs (*i*, *ti*) where *i* is a player in the Bayesian game and *ti* is one of the signals that *i* may receive.
- **Actions:** The set of actions of each player (*i*, *ti*) is the set of actions of player *i* in the Bayesian game.
- **Preferences:** The Bernoulli payoff function of each player (*i*, *ti*) is given by:

$$\sum_{\omega \in \Omega} \Pr(\omega \text{ given } t_i) u_i((a_i, \hat{a}_{-i}(\omega)), \omega)$$

Exercise: An exchange game

- Each of two individuals receives a ticket on which there is an integer from 1 to m indicating the size of a prize she may receive.
- The individuals' tickets are assigned randomly and independently; the probability of an individual's receiving each possible number is positive.
- Each individual is given the option to exchange her prize for the other individual's prize;
- The individuals are given this option simultaneously. If both individuals wish to exchange then the prizes are exchanged; otherwise each individual receives her own prize.

Exercise: An exchange game

- Each individual's objective is to maximize her expected monetary payoff.
- Model this situation as a Bayesian game and show that in any Nash equilibrium the highest prize that either individual is willing to exchange is the smallest possible prize.

Answers

- The following Bayesian game models the situation.
- **Players:** The two individuals.
- **States:** The set of all pairs (s_1, s_2), where *si* is the number on player *i*'s ticket (an integer from 1 to *m*).
- **Actions:** The set of actions of each player is *{Exchange*, *Don't exchange}*.

- **Signals:** The signal function of each player *i* is defined by: $\tau_i(s_1, s_2) = s_i$
(each player observes her own ticket, but not that of the other player).
- **Beliefs:** Type s_i of player *i* assigns the probability $\Pr_j(s_j)$ to the state (s_1, s_2), where *j* is the other player and $\Pr_j(s_j)$ is the probability with which player *j* receives a ticket with the prize s_j on it.

- **Payoffs:** Player *i*'s Bernoulli payoff function is given by
$u_i((X, Y), \omega) = \omega_j$ if $X = Y = Exchange$
and $u_i((X, Y), \omega) = \omega_i$ otherwise.

- Let M_i be the highest type of player i that chooses *Exchange*.
- If $M_i > 1$, then type 1 of player j optimally chooses *Exchange*: by exchanging her ticket, she cannot obtain a smaller prize, and may receive a bigger one.
- Thus if $M_i \geq M_j$ and $M_i > 1$, type M_i of player i optimally chooses *Don't exchange*, because the expected value of the prizes of the types of player j that choose *Exchange* is less than M_i.
- Thus in any possible Nash equilibrium $M_i = M_j = 1$: the only prizes that may be exchanged are the smallest.

Reporting a crime with an unknown number of witnesses

- A crime is observed by a group of 2 people. Each person would like the police to be informed, but prefers that someone else make the phone call.
- Specifically, suppose that each person attaches the value v to the police being informed and bears the cost c if she makes the phone call, where $v > c > 0$.

Reporting a crime with an unknown number of witnesses

- Each of two players does not know whether she is the only witness, or whether there is another witness.
- Denote by π the probability each player assigns to being the sole witness.

Question (a)

- Model this situation as a Bayesian game with three states: one in which player 1 is the only witness, one in which player 2 is the only witness, and one in which both players are witnesses.

Answer Question (a)

- **Players:** The two potential witnesses.
- **States:** "1" (player 1 is the only witness), "2" (player 2 is the only witness),

  and "12" (both players are witnesses).
- **Actions:** Each player's set of actions is {Call,Don't call}.

- **Signals:** Each player receives one of the signals "witness" or "not witness".

  Player 1's signal function τ_1 satisfies:

  $\tau_1(1) = \tau_1(12) =$ "witness"

  and $\tau_1(2) =$ "not witness";

Player 2's signal function τ_2 satisfies:

$\tau_2(2) = \tau_2(12) =$ "witness" and

$\tau_2(1) =$ "not witness".

- **Beliefs:** For i = 1, 2, when player i receives the signal "witness" she assigns probability π to the state "I" and probability $1 - \pi$ to the state "12", and when she receives the signal "not witness", she assigns probability 1 to the state j (where j is the other player).

- **Preferences:** Each player's payoff to an action pair in which at least one player calls is $v - c$ if she calls and v if she does not call; her payoff to the action pair in which neither player calls is 0.

- Note that the concept of a Bayesian game requires us to specify actions for each player independent of the state, so that in this game each player has actions even in the state in which she is not a witness.

- This game is shown in the following Figure, where the action “Call” is denoted “C”, and the action “Don’t call” is denoted “N”.
- In state “1”, only player 1 is a witness, in state “2”, only player 2 is a witness, and in state “12”, both players are witnesses.

π 1 $1 - \pi$ | 1

2 | $1 - \pi$ 2 | π

	C	N
C	*V-C, V*	*V-C, V*
N	0,0	0,0

State ”1”

	C	N
C	*V-C, V-C*	*V-C, V*
N	*V, V-C*	0,0

State ”12”

	C	N
C	*V, V-C*	0,0
N	*V, V-C*	0,0

State ”2”

Question (b)

- Find a condition on π under which the game has a pure Nash equilibrium in which each player chooses *Call* (given the signal that she is a witness).

Answer Question (b)

- A player obtains the payoff $v - c$ if she chooses C and the payoff $(1 - \pi)v$ if she chooses N.
- Thus the game has a pure strategy Nash equilibrium in which each player chooses C in the state in which she is active if and only if

  $v - c \geq (1 - \pi) v$, or $\pi \geq c/v$.
- The action each player chooses in the state in which she is inactive is irrelevant.

Question (c)

- When the condition is violated, find the symmetric mixed strategy Nash equilibrium of the game.

Answer Question (c)

- For a mixed strategy Nash equilibrium in which each player chooses *C* (if she is active) with probability p, where $0 < p < 1$, we need each player's expected payoffs to *C* and *N* to be the same, given that the other player chooses *C* with probability p.

- Thus we need $v - c = (1 - \pi)\, p\, v$, or

$$p = (v - c) / [(1 - \pi)\, v]$$

Answer Question (c)

- If $\pi < c/v$, this number is less than 1, so that the game indeed has a mixed strategy Nash equilibrium in which each player calls with probability p.

Bayes' rule

- The essence of the Bayesian approach is to provide a mathematical rule explaining how you should change your existing beliefs in the light of new evidence.

- In other words, it allows scientists to combine new data with their existing knowledge or expertise.

Bayes' rule

- The canonical example is to imagine that a precocious newborn observes his first sunset, and wonders whether the sun will rise again or not.
- He assigns equal prior probabilities to both possible outcomes, and represents this by placing one white and one black marble into a bag.
- The following day, when the sun rises, the child places another white marble in the bag.
- The probability that a marble plucked randomly from the bag will be white (i.e, the child's degree of belief in future sunrises) has thus gone from a half to two-thirds.

Bayes' rule

- After sunrise the next day, the child adds another white marble, and the probability (and thus the degree of belief) goes from two-thirds to three-quarters.
- And so on.
- Gradually, the initial belief that the sun is just as likely as not to rise each morning is modified to become a near-certainty that the sun will always rise.

Lecture Overview

1. Recap
2. Probability Distributions
3. Conditional Probability
4. Bayes' Theorem

Possible Worlds Semantics

- A random variable is a variable that is randomly assigned one of a number of different values.
- The domain of a variable X, written $dom(X)$, is the set of values X can take.
- A possible world specifies an assignment of one value to each random variable.
- $w \models \phi$ means the proposition ϕ is true in world w.
- Let Ω be the set of all possible worlds.
- Define a nonnegative measure $\mu(w)$ to each world w so that the measures of the possible worlds sum to 1.
- The probability of proposition ϕ is defined by:

$$P(\phi) = \sum_{w \models \phi} \mu(w).$$

Lecture Overview

1. Recap
2. Probability Distributions
3. Conditional Probability
4. Bayes' Theorem

Probability Distributions

Consider the case where possible worlds are simply assignments to one random variable.

Definition (probability distribution)

A probability distribution P on a random variable X is a function $dom(X) \rightarrow [0, 1]$ such that

$$x \mapsto P(X = x).$$

- When $dom(X)$ is infinite we need a probability density function.

Joint Distribution

When there are multiple random variables, their joint distribution is a probability distribution over the variables' Cartesian product

- E.g., $P(X, Y, Z)$ means $P(\langle X, Y, Z \rangle)$.
- Think of a joint distribution over n variables as an n-dimensional table
- Each entry, indexed by $X_1 = x_1, \ldots, X_n = x_n$, corresponds to $P(X_1 = x_1 \wedge \ldots \wedge X_n = x_n)$.
- The sum of entries across the whole table is 1.

Joint Distribution Example

Consider the following example, describing what a given day might be like in Vancouver.

- we have two random variables:
  - $weather$, with domain {Sunny, Cloudy};
  - $temperature$, with domain {Hot, Mild, Cold}.
- Then we have the joint distribution $P(weather, temperature)$ given as follows:

		$temperature$		
		Hot	Mild	Cold
$weather$	Sunny	0.10	0.20	0.10
	Cloudy	0.05	0.35	0.20

Marginalization

Given the joint distribution, we can compute distributions over smaller sets of variables through marginalization:

- E.g., $P(X, Y) = \sum_{z \in dom(Z)} P(X, Y, Z = z)$.
- This corresponds to summing out a dimension in the table.
- The new table still sums to 1.

Marginalization Example

temperature

weather		Hot	Mild	Cold
	Sunny	0.10	0.20	0.10
	Cloudy	0.05	0.35	0.20

If we marginalize out $weather$, we get

$P(temperature) =$

Hot	Mild	Cold
0.15	0.55	0.30

If we marginalize out $temperature$, we get

$P(weather) =$

Sunny	Cloudy
0.40	0.60

Lecture Overview

1. Recap
2. Probability Distributions
3. Conditional Probability
4. Bayes' Theorem

Conditioning

- Probabilistic conditioning specifies how to revise beliefs based on new information.
- You build a probabilistic model taking all background information into account. This gives the prior probability.
- All other information must be conditioned on.
- If evidence e is all of the information obtained subsequently, the conditional probability $P(h|e)$ of h given e is the posterior probability of h.

Semantics of Conditional Probability

- Evidence e rules out possible worlds incompatible with e.
- We can represent this using a new measure, μ_e, over possible worlds

$$\mu_e(\omega) = \begin{cases} \frac{1}{P(e)} \times \mu(\omega) & \text{if } \omega \models e \\ 0 & \text{if } \omega \not\models e \end{cases}$$

Definition

The conditional probability of formula h given evidence e is

$$P(h|e) = \sum_{\omega \models h} \mu_e(w) = \frac{P(h \wedge e)}{P(e)}$$

Conditional Probability Example

		temperature		
		Hot	Mild	Cold
weather	Sunny	0.10	0.20	0.10
	Cloudy	0.05	0.35	0.20

If we condition on $weather = \text{Sunny}$, we get

$$P(temperature|Weather = \text{Sunny}) =$$

Hot	Mild	Cold
0.25	0.50	0.25

Conditioning on $temperature$, we get $P(weather|temperature)$:

		temperature		
		Hot	Mild	Cold
weather	Sunny	0.67	0.36	0.33
	Cloudy	0.33	0.64	0.67

Note that each column now sums to one.

Chain Rule

Definition (Chain Rule)

$$
\begin{aligned}
& P(f_1 \wedge f_2 \wedge \ldots \wedge f_n) \\
&= P(f_n | f_1 \wedge \cdots \wedge f_{n-1}) \times P(f_1 \wedge \cdots \wedge f_{n-1}) \\
&= P(f_n | f_1 \wedge \cdots \wedge f_{n-1}) \times P(f_{n-1} | f_1 \wedge \cdots \wedge f_{n-2}) \times \\
&\quad P(f_1 \wedge \cdots \wedge f_{n-2}) \\
&= P(f_n | f_1 \wedge \cdots \wedge f_{n-1}) \times P(f_{n-1} | f_1 \wedge \cdots \wedge f_{n-2}) \\
&\quad \times \cdots \times P(f_3 | f_1 \wedge f_2) \times P(f_2 | f_1) \times P(f_1) \\
&= \prod_{i=1}^{n} P(f_i | f_1 \wedge \cdots \wedge f_{i-1})
\end{aligned}
$$

E.g., $P(weather, temperature) = P(weather | temperature) \cdot P(temperature)$.

Lecture Overview

Bayes' theorem

The chain rule and commutativity of conjunction ($h \wedge e$ is equivalent to $e \wedge h$) gives us:

$$\begin{aligned} P(h \wedge e) &= P(h|e) \times P(e) \\ &= P(e|h) \times P(h). \end{aligned}$$

If $P(e) \neq 0$, you can divide the right hand sides by $P(e)$, giving us Bayes' theorem.

Definition (Bayes' theorem)

$$P(h|e) = \frac{P(e|h) \times P(h)}{P(e)}.$$

Why is Bayes' theorem interesting?

Often you have causal knowledge:

- $P(symptom \mid disease)$
- $P($*light is off* $\mid$ *status of switches and switch positions*$)$
- $P(alarm \mid fire)$
- $P($*image looks like*  $\mid$ *a tree is in front of a car*$)$

...and you want to do evidential reasoning:

- $P(disease \mid symptom)$
- $P($*status of switches* $\mid$ *light is off and switch positions*$)$
- $P(fire \mid alarm)$.
- $P($*a tree is in front of a car* $\mid$ *image looks like* $)$

Bayes' rule

- Mathematically, Bayes' rule states:
- Posterior = $\dfrac{\text{likelihood * prior}}{\text{marginal likelihood}}$

Bayes' rule

- In this context in which a belief is changed by evidence, the initial belief is called the **prior belief** and the belief modified by the evidence is called the **posterior belief**.

Example

- Suppose you have been tested positive for a disease; what is the probability that you actually have the disease? It depends on the accuracy and sensitivity of the test, and on the background (prior) probability of the disease.

- Suppose that 10% of the population carries the gene *X*, so that in the absence of any other information your prior belief is that you carry the gene with probability 0.1.
- An imperfect test for the presence of *X* is available. The test is positive in 90% of subjects who carry *X* and in 20% of subjects who do not carry *X*.
- The test on you is positive. What should be your posterior belief about your carrying *X*?

- The probabilities are illustrated in the following Figure:

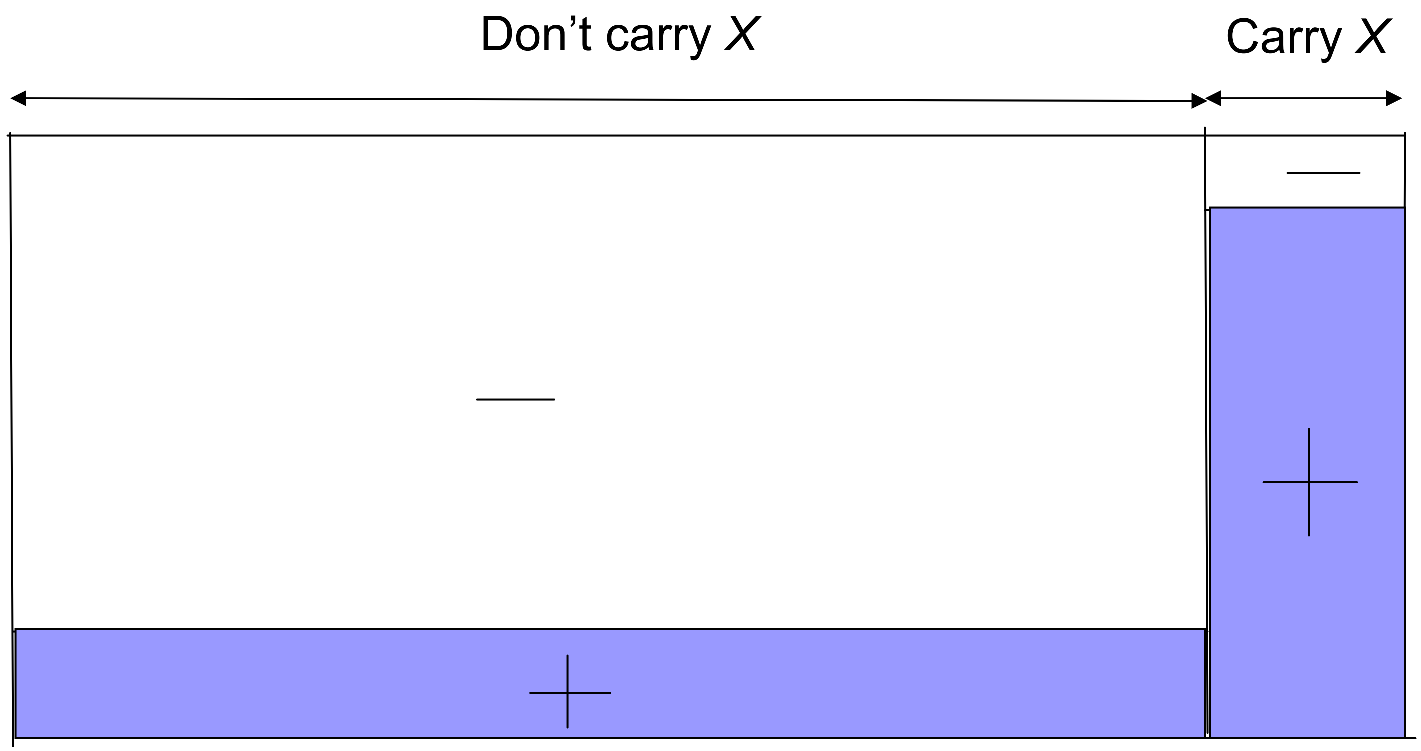

The outer box represents the set of people. People to the right of the vertical line carry gene *X*, while people to the left of this line do not. People in the shaded areas test positive for the gene.

- Consider a random group of 100 people from the population. Of these, on average 10 carry *X* and 90 do not. If all these 100 people were tested, then, on average, 9 of the 10 (90%) who carry *X* and 18 of the 90 (20%) who do not carry *X* would test positive.
- These sets are represented by the shaded areas in the Figure.

- Of all the people who test positive, what fraction of them carry the gene?
- That is, what fraction of the total shaded area in the Figure is the shaded area to the right of the vertical line?
- Of the 100 people, a total of 9 + 18 = 27 test positive, and one-third of these (9/27) carry the gene.
- Thus after testing positive, your **posterior belief** that you carry the gene is 1/3: the positive test raises the probability you assign to your carrying *X* from 1/10 to 1/3.

- To generalize the analysis in this example, we introduce the concept of conditional probability.
- Let *E* and *F* be two events that may be related; assume that Pr(*F*) > 0. Suppose that *F* is true. Define the **probability**
- Pr(*E* | *F*) **of** *E* **conditional on** *F* by:

$$\Pr(E \mid F) = \frac{\Pr(E \ and \ F)}{\Pr(F)}$$

- This number makes sense as the probability that *E* is true *given that F is true*.

- One way to see that it makes sense is to consider the Figure again.
- Let *E* be the event that you carry *X* and let *F* be the event that you test positive.
- If you test positive then we know you lie in the shaded area.
- Given you lie in this area, what is the probability Pr(*E* | *F*) that you lie to the right of the vertical line?
- This probability is the ratio of the shaded area to the right of the vertical line — the probability Pr(*E* and *F*) that you carry the gene and test positive — to the total shaded area—the probability Pr(*F*) that you test positive.

- If the events E and F are independent then

  $\Pr(E \mid F) = \Pr(E)$ and $\Pr(F) > 0$

or, alternatively,

  $\Pr(F \mid E) = \Pr(F)$ and $\Pr(E) > 0$.

- These conditions express directly the idea that the occurrence of one event has no bearing on the occurrence of the other event.

- In using the expression for conditional probability to find the posterior belief in this case, we needed to calculate Pr(*E* and *F*) and Pr(*F*), which were not given directly as data in the problem.
- The data we were given were the prior belief Pr(*E*)=0.1, the probability

  Pr(*F* | *E*)=0.9. of a person who carries the gene testing positive, and the probability Pr(*F* | not *E*)=0.2 of a person who does not carry the gene testing positive.

- Bayes' rule expresses the conditional probability Pr(*E* | *F*) directly in terms of
- Pr(*E*), Pr(*F* | *E*), and Pr(*F* | not *E*):

$$\Pr(E \mid F) = \frac{\Pr(E)\Pr(F \mid E)}{\Pr(E)\Pr(F \mid E) + \Pr(not\ E)\Pr(F \mid not\ E)}$$

- The probability Pr(not *E*) is of course equal to 1 − Pr(*E*); recall that I have assumed that Pr(*F*) > 0.

- So in our case

$$\Pr(E \mid F) = \frac{\Pr(E)\Pr(F \mid E)}{\Pr(E)\Pr(F \mid E) + \Pr(not\ E)\Pr(F \mid not\ E)}$$

$$= \frac{0.1 \times 0.9}{0.1 \times 0.9 + 0.9 \times 0.2} = \frac{9}{27} = \frac{1}{3}$$

Exercise

- Consider a generalization of the example of testing positive for a gene in which the fraction p of the population carry the gene.
- (a) Verify that as p decreases, the posterior probability that you carry X given that you test positive decreases.

Answer Question (a)

- Your posterior probability of carrying *X* given that you test positive is:

$$\frac{\Pr(X)\Pr(positive\ \ test \mid X)}{\Pr(X)\Pr(positive\ \ test \mid X) + \Pr(not\ \ X)\Pr(positive\ \ test \mid not\ \ X)}$$

- We have: Pr(*X*)=*p*, Pr(not *X*)=1-*p*,
  Pr(Positive test | *X*) = 0.9 and Pr(Positive test | not *X*) = 0.2
- This probability is thus equal to:

$$\frac{0.9p}{0.9p + 0.2(1-p)} = \frac{0.9p}{0.2 + 0.7p}$$

This probability is increasing in *p* (i.e. a smaller value of *p* gives a smaller value of the probability).

- (b) What value does this posterior probability take when p is 0.001?

Answer Question (b)

- If $p = 0.001$ then the probability is approximately 0.004.
- That is, if 1 in 1,000 people carry the gene then if you test positive on a test that is 90% accurate for people who carry the gene and 80% accurate for people who do not carry the gene, then you should assign probability 0.004 to your carrying the gene.

- (c) What value does the posterior probability take when p is 0.001 and the test is positive for 99% of those who carry X and is negative for 99% of those who do not carry X?

Answer Question (c)

- If the test is 99% accurate in both cases then the posterior probability is:

$$(0.99 \cdot 0.001) / [0.99 \cdot 0.001 + 0.01 \cdot 0.999] \approx 0.09.$$

Bayes' rule again

- Let us be sure that you understand the Bayes rule.
- A 40-year-old woman undergoes a routine screening for breat cancer, and the test comes back positive. The rate of previously undiagnosed breast cancer for a woman is 0.01, or 1 percent.

- The test is fairly accurate in the sense that if she has cancer, it will produce a positive result 80 percent of the time.
- For a woman who does not have cancer, it will produce a positive result only 10 percent of the time.

Question (a)

- Before the test, what is the probability that the patient has cancer?

Answer to Question (a)

- Since the rate of previously undiagnosed breast cancer for a woman 1 percent, then the probability that this person has cancer is also 1 percent.
- These are here **prior** beliefs.

Question (b)

- What is the probability that the patient *actually* has cancer?

Answer to Question (b)

- Let *E* be the event that a woman has cancer and let *F* be the event that a woman has been tested positive.
- We need to calculate the conditional probability $\Pr(E \mid F)$, that is the probability that a woman has cancer given that the test was positive.

Answer to Question (b)

- According to the Bayes' rule, this conditional probability is equal to:

$$\Pr(E \mid F) = \frac{\Pr(E)\Pr(F \mid E)}{\Pr(E)\Pr(F \mid E) + \Pr(not\ E)\Pr(F \mid not\ E)}$$

- We know that Pr(*E*)=0.01, Pr(*not E*)=0.99.
- Moreover, Pr(*F* | *E*), that is the probability that the test is positive given that a woman has cancer, is 0.80, i.e.

  Pr(*F* | *E*) = 0.80.
- Finally, Pr(*F* | *not E*), that is the probability that the test is positive given that a woman does not have cancer, is 0.10, i.e. Pr(*F* | *not E*) = 0.10.

Answer to Question (b)

- Thus $\Pr(E \mid F) = \dfrac{\Pr(E)\Pr(F \mid E)}{\Pr(E)\Pr(F \mid E) + \Pr(not\ E)\Pr(F \mid not\ E)}$

$$= \frac{0.01 \times 0.8}{0.01 \times 0.8 + 0.99 \times 0.10} = 0.0807$$

- As a result, the probability that this woman has actually cancer is 8.07 percent.
- So the beliefs of this woman having cancer change from 1 percent (before the test) to 8.07 percent (after the test).

Bayes' rule: Fun for a change but more difficult

- Suppose that a cook produces three pancakes, one with one burnt side, one with two burnt sides, one with no burnt sides.
- The cook throws a six-sided die and chooses the pancake at random, which each one having equal probability of the one being served.

- Then the cook flips the pancake high so that each side is equally likely to be the one that shows.
- All you know (besides the way pancake was selected) is that the one on your plate is showing a burnt side on top

Question (a)

- Without the cook and her way of selecting the pancake, what is the probability that you have a pancake with only one burnt side?

Answer to question (a)

- In order to have a pancake with one burnt side, it has to be that it is either a pancake with one burnt side or a pancake with two burnt sides.
- This occurs with probability 1/3 + 1/3 = 2/3 = 0.66.
- These are your prior beliefs.

Question (b)

- Using the information on the cook and her way of selecting pancakes, what is the probability that you have the pancake with only one burnt side?
- Explain first with the counting heuristic and then with Bayes’ rule.

Answer to Question (b)

- **Counting heuristic:**
- Intuitively, the probability that you have the pancake with only one burnt side on your plate has to be:

$$\frac{1}{3}\times 1+\frac{1}{3}\times\frac{1}{2}=\frac{3}{6}=\frac{1}{2}=0.5$$

- Indeed, *either* the pancake with two burnt sides is selected by the cook (this occurs with proba 1/3) and, in that case, with proba 1, you have a pancake on your plate with only a burnt side on top *or* the pancake with one burnt side is selected by the cook (this occurs with proba 1/3) and, in that case, with proba 1/2, you have a pancake on your plate with only a burnt side on top.

Answer to Question (b)

- **Bayes' rule**
- Let *E* be the event that a pancake has one burnt side and let *F* be the event that a pancake has been selected by the cook and end up in your plate.
- We need to calculate the conditional probability Pr(*E* | *F*), that is the probability that you have the pancake with only one burnt side on top.

- According to the Bayes' rule, this conditional probability is equal to:

$$\Pr(E \mid F) = \frac{\Pr(E)\Pr(F \mid E)}{\Pr(E)\Pr(F \mid E) + \Pr(not\ E)\Pr(F \mid not\ E)}$$

- We know that Pr(*E*)=2/3 (prior beliefs), Pr(*not E*)=1/3.
- Moreover, Pr(*F* | *E*), that is the probability that a pancake has been selected by the cook given that it has one burnt side is: 1/3.

- Finally, Pr(*F* | *not E*), that is the probability that you have a pancake with no burnt sides or two burnt sides is 2/3.
- As a result:

$$\Pr(E \mid F) = \frac{\Pr(E)\Pr(F \mid E)}{\Pr(E)\Pr(F \mid E) + \Pr(not\ E)\Pr(F \mid not\ E)}$$

$$= \frac{\frac{2}{3}\frac{1}{3}}{\frac{2}{3}\frac{1}{3} + \frac{1}{3}\frac{2}{3}} = \frac{1}{2}$$

- So you beliefs change from 2/3 (prior) to ½.

Exercise

- Host Monty Hall of the once-popular TV game show ”Let’s Make a Deal” asked his final guest of the day to choose one of three doors (or curtains).
- One door led to the ”grand prize” such as a new car and the other two doors led to ”zonks” or worthless prizes such as goats.

- After the guest chose a door, Monty always opened one of the two other doors to reveal a ”zonk” and always offered the guest the opportunity to switch her choice to the remaining unoponed door.

Question (a)

- What should the guest do if the opened door has been randomly chosen among the two doors non-selected by the guest?

Answer Question (a)

- Let us numbered the doors as follows:
- The door initially chosen by the guest is door number 1, the one opened by Monty is door number 2, and the last door is door number 3.

- At the beginning of the show, the car can be anywhere.
- Denote by V_i the event ”the car is behind the door *i* ”. Thus:

  $Pr(V_1) = Pr(V_2) = Pr(V_3) = 1/3$, these are the prior beliefs.
- Denote by *C* the event ”Monty reveals the zonk”.

- We want to calculate: Pr(V_1 | C), i.e., the probability that the guest chooses the car behind the door given that Monty has opened the door with a zonk. It is:

$$\Pr(V_1 \mid C) = \frac{\Pr(V_1)\Pr(C \mid V_1)}{\Pr(V_1)\Pr(C \mid V_1) + \Pr(not\ V_1)\Pr(C \mid not\ V_1)}$$

- We have: Pr(V_1) = 1/3, Pr(*not* V_1) = 2/3,
- Pr(C | V_1) = 1: This is the proba that Monty has opened the door with a zonk given that the guest has chosen the car behind the door.
- Pr(C | *not* V_1) = ½: This is the proba that Monty has opened the door with a zonk given that the guest has *not* chosen the car behind the door.

- Thus the numerator is equal to 2/3.
- This is very intuitive since it says that the probability that Monty reveals a zonk, is such that:
- Either the guest has chosen the car (this occurs with proba 1/3) and, in that case, Monty automatically reveals a zonk (this occurs with proba 1),
- Or the the guest has *not* chosen the car (this occurs with proba 2/3), and, in that case, Monty will reveal a zonk with proba ½.
- The proba of the numerator, i.e. the proba that Monty reveals a zonk, is thus: (1/3) x 1 + (2/3) x ½ = 2/3.

Answer Question (a)

- As a result, $\Pr(V_1 \mid C) = ½$.
- Thus, once Monty has opened a door and revealed a zonk, the guest is indifferent between the two remaining doors.
- Interestingly, in the real show, very few guests accepted the opportunity to switch, behaving in a Bayesian way.

Question (b)

- What should the guest do if she knows that the door opened by Monty was chosen because it had a zonk behind?

Answer Question (b)

- If the guest knows this information, then $Pr(C \mid not\ V_1)$ is modified.
- We still have: $Pr(V_1) = 1/3$, $Pr(not\ V_1) = 2/3$,
- $Pr(C \mid V_1) = 1$: This is the proba that Monty has opened the door with a zonk given that the guest has chosen the car behind the door.

- **BUT** Pr(C | *not* V_1) = 2/3: This is the proba that Monty has opened the door with a zonk given that the guest has *not* chosen the car behind the door.
- Indeed, to calculate this proba, it has to be that the guest has *not* chosen the car (this occurs with proba 2/3), and, in that case, Monty will reveal a zonk with proba 1 and not ½ as before.

Answer Question (b)

- As a result, $\Pr(V_1 \mid C) = 1/3$, i.e., the probability that the guest chooses the car behind the door given that Monty has opened the door with a zonk.
- Thus, once Monty has opened a door and revealed a zonk, and given that the guest knows that the door opened by Monty was chosen because it had a zonk behind, the guest should definitively switch strategy and chooses the other door since her chance that the car is there is 2/3, i.e. $\Pr(V_3 \mid C) = 2/3$.

- Given that very few guests accepted the opportunity to switch, we are not sure anymore if they were behaving in a Bayesian way. It depends what they know!

Cournot game with imperfect information

Consider the standard Cournot game with two firms competing in quantities. However, one firm does not know the other firm's cost function.

How does the imperfect information affect the firms' behavior?

Assume that both firms can produce the good at constant unit cost. If Q denotes the total quantity produced in the market, then the market price is assumed to be equal to:

$$P(Q) = \begin{cases} \alpha - Q & \text{if } Q \leq \alpha \\ 0 & \text{if } Q > \alpha \end{cases}$$

Assume also that they both know that firm 1's unit cost is c, but only firm 2 knows its own unit cost; firm 1 believes that firm 2's cost is c_L with probability θ and c_H with $1 - \theta$, where $0 < \theta < 1$ and $c_L < c_H$.

Question (a): Model this situation as a Bayesian game.

Players: Firm 1 and firm 2.

States: $\{L, H\}$.

Actions: Each firm's set of actions is the set of its possible outputs (nonnegative numbers).

Signals:

Firm 1's signal function τ_1 satisfies $\tau_1(H) = \tau_1(L)$ (its signal is the same in both states).

Firm 2's signal function τ_2 satisfies $\tau_2(H) \neq \tau_2(L)$ (its signal is perfectly informative of the state).

Beliefs: The single type of firm 1 assigns probability θ to state L and probability $1 - \theta$ to state H. Each type of firm 2 assigns probability 1 to the single state consistent with its signal.

Payoff functions: The firms' Bernoulli payoffs are their profits;

If the actions chosen are (q_1, q_2) and the state is I (either L or H) then firm 1's profit is

$$\pi_1 = q_1 [P(Q) - c] = q_1 (\alpha - c - q_1 - q_2)$$

and firm 2's profit is

$$\pi_2 = q_2 [P(Q) - c_I] = (\alpha - c_I - q_1 - q_2)$$

Question (b): Calculate the Bayesian-Nash equilibrium of this game.

A Nash equilibrium of this game is a triple (q_1^*, q_L^*, q_H^*), where q_1^* is the output of firm 1, q_L^* is the output of type L of firm 2 (i.e. firm 2 when it receives the signal $\tau_2(L)$), and q_H^* is the output of type H of firm 2 (i.e. firm 2 when it receives the signal $\tau_2(H)$), such that:

- q_1^* maximizes firm 1's profit given the output q_L^* of type L of firm 2 and the output q_H^* of type H of firm 2;

- q_L^* maximizes the profit of type L of firm 2 given the output q_1^* of firm 1;

- q_H^* maximizes the profit of type H of firm 2 given the output q_1^* of firm 1.

To find an equilibrium, we first find the firms' best response functions.

Given firm 1's posterior beliefs, its best response $BR_1(q_L, q_H)$ to (q_L, q_L) solves:

$$\max_{q_1} \{\theta\,[P(q_1+q_L)-c]\,q_1+(1-\theta)\,[P(q_1+q_H)-c]\,q_1\}$$

$$\Leftrightarrow \max_{q_1} \{\theta\,(\alpha-c-q_1-q_L)\,q_1+(1-\theta)\,(\alpha-c-q_1-q_H)\,q_1\}$$

First-order condition gives:

$$\theta\,(\alpha-c-2q_1-q_L)+(1-\theta)\,(\alpha-c-2q_1-q_H)=0$$

which leads to the following best-reply function:

$$BR_1(q_L, q_H) \equiv q_1(q_L, q_H) = \frac{(\alpha-c)}{2} - \frac{[\theta q_L+(1-\theta)\,q_H]}{2}$$

As a result:

$$BR_1(q_L, q_H) = \begin{cases} \frac{(\alpha - c)}{2} - \frac{[\theta q_L + (1-\theta) q_H]}{2} & \text{if } \theta q_L + (1-\theta)\, q_H \leq \alpha - c \\ 0 & \text{otherwise} \end{cases}$$

Firm 2's best response $BR_L(q_1)$ to q_1 when its cost is c_L solves:

$$\max_{q_L} \{[P(q_1 + q_L) - c_L]\, q_L\}$$

$$\Leftrightarrow \max_{q_L} \{(\alpha - c_L - q_1 - q_L)\, q_L\}$$

The first-order condition is:

$$(\alpha - c_L) - q_1 - 2q_L = 0$$

which leads to the following best-reply function:

$$BR_L(q_1) \equiv q_L(q_1) = \frac{(\alpha - c_L)}{2} - \frac{q_1}{2}$$

As a result:

$$BR_L(q_1) = \begin{cases} \frac{(\alpha - c_L)}{2} - \frac{q_1}{2} & \text{if } q_1 \leq \alpha - c \\ 0 & \text{otherwise} \end{cases}$$

Firm 2's best response $BR_H(q_1)$ to q_1 when its cost is c_H solves:

$$\max_{q_H} \{[P(q_1 + q_H) - c_H]\, q_H\}$$

which is equivalent to:

$$\max_{q_H} \{(\alpha - c_H - q_1 - q_H)\, q_H\}$$

The first-order condition is:

$$(\alpha - c_H) - q_1 - 2q_H = 0$$

which leads to the following best-reply function:

$$BR_H(q_1) \equiv q_H(q_1) = \frac{(\alpha - c_H)}{2} - \frac{q_1}{2}$$

As a result:

$$BR_H(q_1) = \begin{cases} \frac{(\alpha - c_H)}{2} - \frac{q_1}{2} & \text{if } q_1 \leq \alpha - c \\ 0 & \text{otherwise} \end{cases}$$

A Nash equilibrium is a triple (q_1^*, q_L^*, q_H^*) such that

$$q_1^* = BR_1(q_L^*, q_H^*)$$

$$q_L^* = BR_L(q_1^*)$$

$$q_H^* = BR_H(q_1^*)$$

Solving these equations under the assumption that they have a solution in which

all three outputs are positive, we obtain

$$q_1^* = \frac{(\alpha - 2c)}{3} + \frac{[\theta q_L + (1-\theta) q_H]}{3}$$

$$q_L^* = \frac{(\alpha - 2c_L + c)}{3} - \frac{(1-\theta)(c_H - c_L)}{6}$$

$$q_1^* = \frac{(\alpha - 2c_H + c)}{3} + \frac{\theta (c_H - c_L)}{6}$$

It is easy to verify that for values of c_H and c_L close enough, this Bayesian-Nash equilibrium is such that all outputs are positive.

Relation between mixed-strategy equilibrium and Bayesian equilibrium (Gibbons)

A man (Patrick) and a woman (Lindsey) are trying to decide on an evening's entertainment. The following normal form game (Battle of Sexes) represents their payoffs.

		Patrick (Player 2)	
		Opera	Fight
Lindsey	Opera	$(2, 1)$	$(0, 0)$
(Player 1)	Fight	$(0, 0)$	$(1, 2)$

By convention, the payoff to row-player (here Lindsey, player 1) is the first payoff given, followed by the payoff of the column player (here Patrick, player 2).

Question (a): Find the mixed-strategy equilibria of this game.

We have seen that this game has three mixed strategy Nash equilibria:

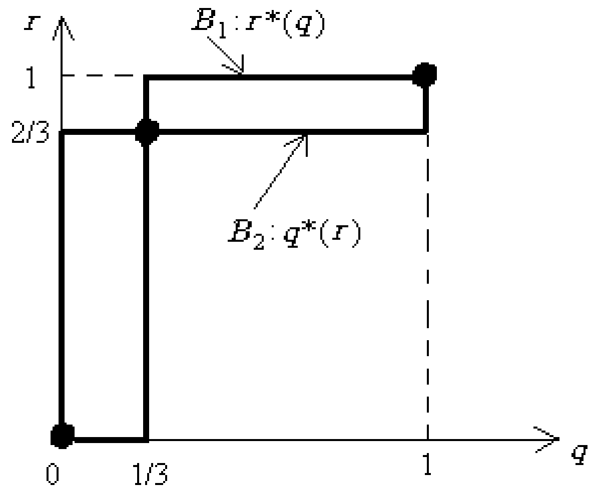

(i) $(q, 1-q) = (0, 1)$ for Patrick and $(r, 1-r) = (0, 1)$ for Lindsey, i.e. the pure strategy equilibrium (Fight, Fight)

(ii) $(q, 1-q) = (1, 0)$ for Patrick and $(r, 1-r) = (1, 0)$ for Lindsey, i.e. the pure strategy equilibrium (Opera, Opera),

(iii) the mixed strategies $(q, 1-q) = (1/3, 2/3)$ for Patrick and $(r, 1-r) = (2/3, 1/3)$ for Lindsey are a Nash equilibrium.

Question (b): We still suppose that Patrick and Lindsey are trying to decide on an evening's entertainment but now Lindsey is unsure whether Patrick prefers to go out with her or prefers to avoid her and Patrick is also unsure whether Lindsey prefers to go out with him or prefers to avoid him. The following normal form of the game gives their payoffs.

		Patrick (Player 2)	
		Opera	Fight
Lindsey	Opera	$(2+t_L, 1)$	$(0,0)$
(Player 1)	Fight	$(0,0)$	$(1, 2+t_P)$

We assume that t_L is privately known by Lindsey and t_P is privately known by Patrick. We assume that t_L and t_P are independent draws from a uniform distribution on $[0, x]$.

Give a formal definition of this static Bayesian game in normal form (action spaces, type spaces, beliefs).

Consider the following Bayesian game:

$$G = \{A_L, A_P; T_L, T_P; p_L, p_P; u_L, u_P\}$$

Players: Lindsey and Patrick;

Strategies: $A_L = A_P = \{\text{Opera , Fight}\}$;

Type spaces: $T_L = T_P = [0, x]$;

Beliefs: $p_L(t_P) = p_P(t_L) = 1/x$ for all t_L and t_P;

Payoffs: there are given the matrix above.

Question (c): We now have the following strategies.

Lindsey plays "Opera" if $t_L \geq c$ (where c is some critical value) and plays "Fight" otherwise. Similarly, Patrick plays "Fight" if $t_P \geq p$ (where p is some critical value) and plays "Opera" otherwise. Observe that $0 < c < 1$ and $0 < p < 1$.

When the distribution is uniform, what is the probability that Lindsey plays "Opera" and the probability that Patrick plays "Fight"? What are these probabilities when the distribution is not uniform?

When the distribution is uniform, the probability that Lindsey plays "Opera" is $(x - c)/x$ and the probability that Patrick plays "Fight" is $(x - p)/x$.

If the distribution was not uniform, then the probability that Lindsey plays "Opera" would be $1 - F(c)$ since $F(c) = \Pr(t_L \leq c)$ and the probability that Patrick plays "Fight" would be $1 - F(p)$ since $F(p) = \Pr(t_P \leq p)$.

Question (d): Let us assume again that t_L and t_P are independent draws from a *uniform* distribution on $[0, x]$. Then, if Lindsey and Patrick are playing the strategies describe above, what are the values of c and p such that these strategies are a Bayesian Nash equilibrium?

Given Patrick's strategy, Lindsey's expected payoff from playing "Opera" is

$$\frac{p}{x}(2 + t_L) + \left[1 - \frac{p}{x}\right].0 = \frac{p}{x}(2 + t_L)$$

and Lindsey's expected payoff from playing "Fight" is

$$\frac{p}{x}.0 + \left[1 - \frac{p}{x}\right].1 = 1 - \frac{p}{x}$$

Thus, for Lindsey, playing "Opera" is optimal if and only if:

$$t_L \geq \frac{x}{p} - 3 \equiv c \qquad (1)$$

Similarly, given Lindsey's strategy, Patrick's expected payoff from playing "Fight" is

$$\left[1 - \frac{c}{x}\right] \times 0 + \frac{c}{x}(2 + t_P) = \frac{c}{x}(2 + t_P)$$

and Patrick's expected payoff from playing "Opera" is

$$\left[1 - \frac{c}{x}\right] \times 1 + \frac{c}{x} \times 0 = 1 - \frac{c}{x}$$

Thus, for Patrick, playing “Fight” is optimal if and only if:

$$t_P \geq \frac{x}{c} - 3 \equiv p \tag{2}$$

Solving (1) and (2) simultaneously yields

$$\begin{cases} x - 3p = cp \\ x - 3c = cp \end{cases}$$

or equivalently

$$\begin{cases} p = c \\ p^2 + 3p - x = 0 \end{cases}$$

The discriminant of the quadratic equation $p^2 + 3p - x = 0$ is:

$$\Delta = 9 + 4x$$

and the two roots are given by

$$p_1 = \frac{-3 - (9 + 4x)^{1/2}}{2} < 0$$

and

$$p_2 = \frac{-3 + (9 + 4x)^{1/2}}{2} > 0 \qquad (3)$$

Since p has to be positive, then only p_2 is possible. Thus the probability that Lindsey plays "Opera", namely $(x-c)/x$, and the probability that Patrick plays "Fight", namely $(x-p)/x$, are both equal to

$$\frac{x-c}{x} = \frac{x-p}{x} = 1 - \frac{-3+(9+4x)^{1/2}}{2x}$$

To summarize, the values of c and p such the strategies describe above (i.e. Lindsey plays "Opera" if $t_L \geq c$ and plays "Fight" otherwise; Patrick plays "Fight" if $t_P \geq p$ and plays "Opera" otherwise) are a Bayesian Nash equilibrium are

$$c = p = \frac{-3+(9+4x)^{1/2}}{2}$$

Question (e): Show that when x approaches zero, i.e. the incomplete information disappears, the players' behavior in this pure-strategy Bayesian Nash equilibrium approaches their behavior in the mixed-strategy Nash equilibrium in the original game of complete information (question 1a).

We have to show that both $(x-c)/x$ and $(x-p)/x$ approach 2/3 as x approaches 0. It is easy to see that

$$\lim_{x \to 0} \frac{x-c}{x} = \lim_{x \to 0} \frac{x-p}{x} = \lim_{x \to 0} \left[1 - \frac{-3 + (9+4x)^{1/2}}{2x} \right]$$

is not determined. We have to use the l'Hospital rule. The latter says:

If two functions $f()$ and $g()$ are continuous and differentiable, then

$$\lim_{x \to 0} \frac{f(x)}{g(x)} = \lim_{x \to 0} \frac{f'(x)}{g'(x)}$$

Denote by

$$f(x) = -3 + (9 + 4x)^{1/2}$$

and

$$g(x) = 2x$$

Then, since

$$f'(x) = 2/\left(9 + 4x\right)^{1/2}$$

and

$$g'(x) = 2$$

we have:

$$
\begin{aligned}
& \lim_{x \to 0} \left[\frac{-3 + (9+4x)^{1/2}}{2x} \right] \\
= & \lim_{x \to 0} \left[\frac{1}{(9+4x)^{1/2}} \right] = \frac{1}{3}
\end{aligned}
$$

As a result,

$$
\begin{aligned}
\lim_{x \to 0} \frac{x-c}{x} &= \lim_{x \to 0} \frac{x-p}{x} \\
&= \lim_{x \to 0} \left[1 - \frac{-3 + (9+4x)^{1/2}}{2x} \right] \\
&= 1 - \frac{1}{3} = \frac{2}{3}
\end{aligned}
$$

Thus, as the incomplete information disappears, the players' behavior in the pure-strategy Bayesian Nash equilibrium of the incomplete-information game approaches their behavior in the mixed-strategy Nash equilibrium in the original game of complete information.

This shows that a mixed-strategy Nash equilibrium in a game of complete information can (almost always) be interpreted as a pure-strategy Bayesian Nash equilibrium in a closely related game with a little bit of incomplete information.