Game Theory with Applications to Finance and Marketing

Some Examples

Instructor: Chyi-Mei Chen
Room 1102, Management Building 2
(02) 3366-1086
(email) cchen@ccms.ntu.edu.tw

1. **Definition 1.** A game described by *(i) the set of players, (ii) the strategies available to each player, and (iii) the payoff of each player as a function of the vector of all players' strategic choices* is called a game depicted in *normal form.*

2. **Example 1**. The following is a two-player normal-form game.

player 1/player 2	L	R
U	0,1	-1,2
D	2,-1	-2,-2

- Who are the players? Players 1 and 2.
- What strategies are available to player 1? U and D. What strategies are available to player 2? L and R.
- What does each player get given players' choices of strategies? Players 1 and 2 get respectively 0 and 1, if the vector of the two players' strategies is (U, L); that is, if player 1 plays U and player 2 plays L. Let us write

$$u_1(U, L) = 0,\ u_2(U, L) = 1. \tag{1}$$

Similarly, we have

$$u_1(U, R) = -1,\ u_2(U, R) = 2,\ u_1(D, L) = 2, \tag{2}$$

$$u_2(D, L) = -1,\ u_1(D, R) = -2,\ u_2(D, R) = -2. \tag{3}$$

The functions $u_1(\cdot, \cdot)$ and $u_2(\cdot, \cdot)$ are referred to as the two players' *payoff functions.*

3. **Definition 2**. A (pure strategy) Nash equilibrium (NE) for a two-player normal-form game where the set of player 1's strategies is X (hereafter referred to as player 1's *strategy space*) and the set of player 2's strategies is Y (hereafter referred to as player 2's *strategy space*) is a pair (x^*, y^*) such that $x^* \in X$, $y^* \in Y$, and

$$u_1(x^*, y^*) \geq u_1(x, y^*), \quad \forall x \in X, \tag{4}$$

$$u_2(x^*, y^*) \geq u_2(x^*, y), \quad \forall y \in Y. \tag{5}$$

The game in Example 1 has $X = \{U, D\}$ and $Y = \{L, R\}$. This game has two pure-strategy NEs, (U,R) and (D,L).

4. **Definition 3**. A *mixed strategy* for player 1 in the above two-player normal-form game is a probability distribution f over X, and a mixed strategy for player 2 is a probability distribution g over Y. A *mixed strategy Nash equilibrium* is a pair (f, g) such that facing g, f is one of player 1's best choices over X, and facing f, g is one of player 2's best choices over Y.

   Recall the game in Example 1. The probability distribution $f(U) = f(D) = \frac{1}{2}$ is one mixed strategy for player 1. The probability distribution $g(R) = \frac{1}{3}$ $g(L) = \frac{2}{3}$ is a mixed strategy for player 2. It turns out that this pair (f, g) is the unique (non-degenerate) mixed-strategy NE for the game in Example 1, because

$$g(L)u_1(U, L) + g(R)u_1(U, R) = g(L)u_1(D, L) + g(R)u_1(D, R), \tag{6}$$

$$f(U)u_2(U, L) + f(D)u_2(D, L) = f(U)u_2(U, R) + f(D)u_2(D, R). \tag{7}$$

   There is an obvious reason for the above two equations: if given his rival's mixed strategy, a player strictly prefers one pure strategy to the other, then he will assign zero probability to the latter; that is, a mixed strategy can never be his best response. Thus in a mixed strategy Nash equilibrium, where each player assigns a positive probability to every pure strategy, a player has to feel *indifferent* about his two pure strategies.

5. **Example 2**. Consider a two-player game, where the two must pick an integer from the set $\{1, 2, \cdots, 100\}$ at the same time. If they pick the same number, then they each get 1; or else, they each get zero. Find the pure strategy NE's. Find the mixed strategy NE's.

6. **Definition 4**: A dynamic game is usually described in extensive form and represented by a *game tree*. For example, consider the following game tree:

$$(1)\text{—}\left[\begin{array}{ccc} Up\text{—} & (2)\text{—} & \left[\begin{array}{cc} Right\text{—} & (0,1) \\ Left\text{—} & (-1,2)\end{array}\right. \\ & \vdots & \\ Down\text{—} & (2)\text{—} & \left[\begin{array}{cc} Right\text{—} & (2,-1) \\ Left\text{—} & (-2,-2)\end{array}\right.\end{array}\right.$$

   In this game tree, the first mover's (player 1's) decision node is the root of the tree, and each of the two pure strategies available to the first mover is represented by a branch emanating from that decision node. These branches, labeled up and down respectively, lead to the second mover's decision nodes. Note that player 2's two decision nodes are connected by a dotted line, and these two decision nodes define player 2's *information set* at the time player 2 must choose between right and left. Formally, an *information set* is a set of decision nodes for a player, who, while knowing that he is sitting on one of those nodes contained in the information set, cannot tell which node in the information set he is exactly sitting on.The remaining game tree starting from a *singleton* information set is called a *subgame* of the original game.

7. **Definition 5**. A subgame perfect Nash equilibrium (SPNE) is an NE for a game described as a game tree, which specifies NE strategies in each and every subgame.

8. **Example 3**. A bricks-and-mortar store F can produce and sell product X costlessly to 2 consumers A and B, where A knows everything about internet, while B has no knowledge about it. F cannot distinguish A from B, but F knows that A (respectively, B) is willing to pay

2 (respectively, 5) dollars for 1 unit of X, and they both have unit demand for X.
(i) First suppose that e-commerce is unavailable. F first announces price p, and given p, A and B decide whether to buy X from F. What is p in equilibrium? What is F's profit?
(ii) Now suppose instead that F can first spend $t > 0$ and set up an online outlet, and if t is spent, then F would announce the online price q and the offline price p to A and B. A can then decide where (at the online or offline outlet) to buy X, but B can only buy from F's offline store. Should F spend t, and what is F's equilibrium prices?
(iii) Re-consider (ii) but assuming that B's reservation value is 3.9.
(iv) Would a higher demand would make it more likely that t is spent?

Solution. For part (i), F should set p at 5 dollars, which is also F's profit. For part (ii), if t is spent, then F should announce $p = 5$ and $q = 2$, so that F would gain $q - t$ by spending t. Thus F should spend t if and only if $t < 2$.[1] For part (iii), F's profit is 4 dollars if giving up the online outlet, while selling through both online and offline outlets would yield a profit of $3.9 + 2 - t$, and hence F should sell through both channels if $t < 1.9$.[2]

Note that the assumption that B has a higher reservation value than A does is crucial in the above analysis. If the reservation values are reversed, F cannot price discriminate between A and B.

Finally, for part (iv), note that F's incentive to spend t may be reduced when A's reservation value increases: if A is already served without the internet, then the benefit from spending t is equal to the difference in

[1]Note that A is not served without on-line markets, and t is the cost that F incurs in order to extract the 2 dollars from A in the presence of on-line markets.

[2]Note that in (iii) selling to A online would create a loss: F would charge A slightly less than 2 dollars, but F has to spend $t > 0$. Spending t can still be beneficial because F can then charge B 3.9 instead of 2 dollars. Why? With A being directed to the on-line markets, B is identified as the only segment left to be served at the original store, and hence F can fully extract the consumer surplus from B. The extra 1.9=3.9-2 dollars that F can make from serving B must be greater than the loss incurred when F moves A from the offline outlet to the online outlet. Thus F should take this *dual channels* strategy if and only if $t < 1.9$.

A's and B's reservation values, and this difference *decreases* with A's reservation value.

For example, let $t = 1.85$ and B's valuation for X be 3.9. Suppose that A's valuation for X now increases from $v_A = 2$ to $v_A = 2.1$. If $v_A = 2$, then F would get $3.9 - 2 - t > 0$ by spending t; but F would get $3.9 - 2.1 - t < 0$ instead, if $v_A = 2.1$.

On the other hand, an increase in B's demand does weakly encourage F to spend t: it does not affect F's preference about spending t if F would choose not to serve A in the absence of the internet, but it weakly increases F's benefit from spending t if F would choose to serve A in the absence of the internet.

9. We have thus far assumed that each player knows the rivals' payoff functions. Such a game is a game with *complete (or symmetric) information*. What if at least one player in the game does not know for sure another player's payoff function? We call it a game with *information asymmetry*, or a game with *incomplete information*, or simply a *Bayesian game*. In a Bayesian game, at least one player Z has more than one possible payoff function. We say that this player Z has more than one *type*. At least one other player W cannot be sure which type player Z has. In this case, we shall look for an equilibrium called *Bayesian equilibrium* (BE). This is nothing but a Nash equilibrium of an enlarged version of the original game, where each different type of Z is now treated as a distinct player.

10. **Definition 6**: Recall the two-player game where player 1's strategy space is X and player 2's strategy space is Y, and assume now that player 2 has two possible types (or two possible payoff functions), θ_1 and θ_2, which, from player 1's perspective, may occur with probabilities π_1 and π_2 respectively. Certainly, player 2 knows his own type for sure.

    For all $x \in X$, $y \in Y$, and $\theta \in \{\theta_1, \theta_2\}$, let $u_1(x, y)$ be player 1's payoff, and $u_2(x, y; \theta)$ the type-θ player 2's payoff. (This is referred to as a private-value model. If u_1 also depends on θ, then this is a common-value model.) A Bayesian equilibrium for this two-player game is nothing but the Nash equilibrium of the three-player game where the two types of player 2 are treated as two different players. Thus a Bayesian equilibrium is a triple (x^*, y_1^*, y_2^*) such that $x^* \in X$, $y_1^* \in Y$, $y_2^* \in Y$, and the following three *incentive compatibility* conditions hold:

$$\pi_1 u_1(x^*, y_1^*) + \pi_2 u_1(x^*, y_2^*) \geq \pi_1 u_1(x, y_1^*) + \pi_2 u_1(x, y_2^*), \quad \forall x \in X; \quad (8)$$

$$u_2(x^*, y_1^*; \theta_1) \geq u_2(x^*, y; \theta_1), \quad \forall y \in Y; \quad (9)$$

$$u_2(x^*, y_2^*; \theta_2) \geq u_2(x^*, y; \theta_2), \quad \forall y \in Y. \quad (10)$$

In words, x^* is player 1's best response, which is *on average* the optimal strategic choice of player 1. It is not really player 1's best response against player 2 if player 1 is sure that player 2 will use y_1^*. Neither is it player 1's best response against player 2 if player 1 is sure that player 2 will use y_2^*. Since player 1 can only choose one x in X to play against two possible types of player 2, given his conjecture of (y_1^*, y_2^*), the choice x^* must be *on average optimal*. On the other hand, player 2 knows his own type, and his best response against player 1's average optimal choice x^* depends on his type. Note that θ denotes player 2's type, and it determines $u_2(x, y)$! This is why we say that incomplete information in this game is equivalent to player 1 not knowing player 2's payoff function. Again, player 1's average optimal choice x^*, player 2's best response y_1^* when his type is θ_1, and player 2's best response y_2^* when his type is θ_2, must altogether form a Nash equilibrium. This three-player Nash equilibrium is what we defined as the *Bayesain equilibrium*.

11. **Example 4**. In a Cournot duopoly with a homogeneous product, firms 1 and 2 must simultaneously choose supply quantities q_1 and q_2, and given q_1 and q_2, the equilibrium product price would be

$$P(q_1 + q_2) = \tilde{a} - q_1 - q_2, \quad (11)$$

where the random variable $\tilde{a}$ may take on 2 with probability $\frac{1}{3}$ or 4 with probability $\frac{2}{3}$. Firm 1 knows the realization of $\tilde{a}$ when choosing q_1, but firm 2 only knows the distribution of $\tilde{a}$. Find a BE, assuming that firms seek to maximize expected profits.

Solution. First observe that firm 1 has two possible types. So, we shall consider a three-player game, where the two types of firm 1 will be regarded as two different players, and we shall look for the NE of this 3-player game.

By definition, we must find three supply quantities $\{q_1^*(2), , q_1^*(4), q_2^*\}$, such that, given any two of them, the third one is the corresponding player's best response

In other words, for firm 2,

$$q_2^* = arg \max_{q_2} \frac{1}{3}[q_2(2 - q_1^*(2) - q_2)] + \frac{2}{3}[q_2(4 - q_1^*(4) - q_2)].$$

Similarly, for the firm 1 that has seen $\tilde{a} = 2$,

$$q_1^*(2) = arg \max_{q_1} q_1(2 - q_1 - q_2^*);$$

and for the firm 1 that has seen $\tilde{a} = 4$,

$$q_1^*(4) = arg \max_{q_1} q_1(4 - q_1 - q_2^*).$$

Solving the above system of equations, we obtain the NE of the three-player game, or the BE of the original game, which is

$$(q_1^*(2), q_1^*(4), q_2^*) = (\frac{4}{9}, \frac{13}{9}, \frac{10}{9}).$$

12. **Example 5**. Consider firms 1 and 2 competing in the following Cournot game. The inverse demand is

$$p = \tilde{a} - q_1 - q_2,$$

where $\tilde{a}$ is equally likely to take on 4 or 2. Firms can operate costlessly.

First suppose that the firms compete after seeing the realization of $\tilde{a}$. In this case, they choose

$$q_1 = q_2 = \frac{\tilde{a}}{3}$$

in state $\tilde{a}$, and each firm has equilibrium payoff equal to

$$\frac{1}{2}[\frac{16}{9} + \frac{4}{9}] = \frac{10}{9}.$$

Next suppose that the firms must compete before knowing the realization of $\tilde{a}$. In this case, given q_j, firm i seeks to

$$\max_{q_i} q_i(E[\tilde{a}] - q_i - q_j).$$

Consequently, they choose

$$q_1 = q_2 = \frac{E[\tilde{a}]}{3} = 1.$$

Each firm gets the payoff

$$(2 - 1 - 1) \cdot 1 = 0$$

in state $\tilde{a} = 2$, and the payoff

$$(4 - 1 - 1) \cdot 1 = 2$$

in state $\tilde{a} = 4$. Their common expected profit is $1 < \frac{10}{9}$.

Finally, assume that before competing firm 1 knows the realization of $\tilde{a}$ but firm 2 does not. In this case, their output choices in the Bayesian equilibrium are such that

$$q_2 = \frac{6 - q_1(4) - q_1(2)}{4}, \; q_1(4) = \frac{4 - q_2}{2}, \; q_1(2) = \frac{2 - q_2}{2},$$

so that

$$q_1(4) = \frac{3}{2} > \frac{4}{3} > q_2 = 1 > \frac{2}{3} > q_1(2) = \frac{1}{2}.$$

Firm 1's expected profit is

$$\frac{1}{2}[4 - \frac{3}{2} - 1] \cdot \frac{3}{2} + \frac{1}{2}[2 - \frac{1}{2} - 1] \cdot \frac{1}{2} = \frac{5}{4} > \frac{10}{9}.$$

Firm 2's expected profit is

$$\frac{1}{2}[4 - \frac{3}{2} - 1] \cdot 1 + \frac{1}{2}[2 - \frac{1}{2} - 1] \cdot 1 = 1 < \frac{10}{9}.$$

Thus a firm does benefit (suffer) from its superior (inferior) demand information.

Nonetheless, firm 1 can do better by sharing information with firm 2 in state $\tilde{a} = 2$: being uninformed firm 2 is producing too much in this low-demand state, which forces firm 1 to cut back its production, and letting firm 2 know the demand state will benefit firm 1.

However, if firm 2 is rational, then firm 2 knows that firm 1 is willing to share the demand information with it if and only if the demand state is $\tilde{a} = 2$! Thus we have a separating equilibrium where firm 1 would lose its high payoff due to superior information (recall that its payoff with information asymmetry is $\frac{5}{4} > \frac{10}{9}$!).

Firm 1 would be better off to "commit" to never sharing demand information with firm 2.

13. Now we give a formal definition of Bayesian game.
Definition 7. Given a game, an event is *mutual knowledge* if every player knows it. Given a game, an event is called the players' *common knowledge* if every player knows it, everyone knows that everyone knows it, everyone knows that everyone knows that everyone knows it, and so on. Anything that is not common knowledge is some player's *private information.* A player's private information is also called his *type.* A game where no players have private information (everything relevant is common knowledge) is *a game with complete information.* Otherwise, the game is one with *incomplete information*, or one with *information asymmetry.*

14. A village has three residents, and they all know that a resident's hair can be either red (R) or black (B). A resident can see the color of each neighbor's hair, but does not know the color of his own hair. The three residents are not allowed to communicate in any way. They must meet (quietly) for 1 hour at 9am each day, trying to figure out the color of his own hair. Suppose that a resident that figured out the color of his own hair by the evening of date t would be allowed (by Trump, unlikely?) to immigrate to the USA in the evening of date t. The three residents all wish to immigrate to the USA as early as they can. Suppose that exactly one of the three resident has black hair.
(i) Suppose that the three residents' first meeting is at date 1. When would a red head get to immigrate to the USA?
(ii) Suppose that at date $n \geq 1$, an honest person passed through the

village at 9:20am (whose honesty is well known to the residents), and he told the three residents during their daily meeting that *at least one of them has red hair*. At which date would a red head get to immigrate to the USA? At which date would the black head get to immigrate to the USA?

15. **Definition 8.** An incomplete-information game where at least one uninformed player can act after observing an informed player's action is called a *dynamic game*. A dynamic Bayesian game is called a *signaling game* if there are only two players, one informed and the other uninformed, each having one move, and the uninformed moves right after seeing the informed's move, with the game ending right after the uninformed makes his move.

16. **Example 6**. Consider the following TV game show, where three boxes are presented to a guest (G) by the host (H). G understands that H knows which of the three boxes contains a prize even before the show begins. The show proceeds in 3 steps as follows. (Step 1.) G would have to choose one box. (Step 2.) Then H would open another box for G, and if the opened box contains the prize, then the prize is given to G; or else, (Step 3.) G can choose to or not to swap the box that G chose in Step 1 with the box that neither G nor H has touched.

    Now, suppose that the show has finished Step 2, and H did not open a box containing the prize. Should G make the swap in Step 3?

    (i) First assume that H would like to give the prize to G whenever possible (type a).
    (ii) Next assume instead that H would prevent G from getting the prize whenever possible (type b).
    (iii) Now, suppose that G believes that H may be type a for probability α. Then G should make the swap if and only if $\alpha \leq \alpha^*$, where $\alpha^* =$?