

2. Simultaneous-Move Games

We now want to study the central question of game theory: how should a game be played. That is, what should we expect about the strategies that will be played in a game. We will start from the simplest types of games, those of simultaneous-move games in which all players move only once and at the same time. We will first study simultaneous games with complete information, and then study simultaneous games with incomplete information, where each player's payoffs may be known only by the player.

Dominant strategy and dominated strategy

One quite convincing way of reaching a prediction in a game is the idea of dominance. Let's start from a simple example.

The Prisoner's dilemma. Two individuals are arrested for allegedly committing a crime. The district attorney (DA) does not have enough evidence to convict them. The two suspects are put in two separate jail cells, and are told the following: if one confesses and another does not confess, then the confessed suspect will be rewarded with a light sentence of 1 year, while the unconfessed prisoner will be sentenced to 10 years; if both confess, then both will be sentenced to 5 years; if neither confesses, then the DA will still have enough evidence to convict them for a less serious crime, and sentence each of them to jail for 2 years. The game is depicted below:

| | | |
|------------|------------|---------|
| | Prisoner 2 | |
| | DC | C |
| DC | -2,-2 | -10, -1 |
| Prisoner 1 | | |
| C | -1, -10 | -5, -5 |

Each player has two possible strategies: DC or C. What strategy will each player choose? Notice that strategy C is best for each player regardless of what the other

player's strategy is. In this case, we should expect that each player will choose C. In this game, C is said to be dominant strategy for each player. That is, a palyer's dominant strategy is a strategy that is best for the player regardless of what other players do. The game is called the prisoner's dilemma because both players could have achieved higher payoffs if they both choose DC, but this outcome will not be achieved when each player seeks to maximize his own payoff.

Definition. A strategy $s_i \in S_i$ is a strictly dominant strategy for player i in game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if for all $s'_i \neq s_i$, we have

$$u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$$

for all $s_{-i} \in S_{-i}$.

A related concept is dominated strategy. Consider the following game:

| | | | |
|----------|---|----------|--------|
| | | Player 2 | |
| | | L | R |
| Player 1 | U | 1,-1 | -1, 1 |
| | M | -1, 1 | 1, -1 |
| | D | -2, 5 | -3, -2 |

Neither player in the game has a dominant strategy. But no matter what 2 does, 1 does better playing U (or M) than playing D. Thus a rational player should not play D. D in this case is called a strictly dominated strategy. Formally,

Definition. A strategy $s_i \in S_i$ is called a strictly dominated strategy for player i in game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if there is another strategy $s'_i \in S_i$ such that

$$u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i})$$

for all $s_{-i} \in S_{-i}$. In this case we say that strategy s'_i strictly dominates s_i .

We should expect a player not to play a strictly dominated strategy.

Thus, a strictly dominant strategy for player i is a strategy that strictly dominates all other player i's strategies.

Still another relevant concept is weakly dominated strategy. Consider the following game:

| | | Player 2 | |
|----------|---|----------|------|
| | | L | R |
| Player 1 | U | 5, 1 | 4, 0 |
| | M | 6, 0 | 3, 1 |
| | D | 6, 4 | 4, 4 |

1's payoff from D is at least as high as from M, whether 2 chooses L or R, and is strictly higher if 2 chooses R. In this case, we say that M is a weakly dominated strategy for player 1. In this case, we say that M is a weakly dominated strategy for player 1. Similarly, 1's payoff from D is at least as high as from U, whether 2 chooses L or R, and is strictly higher if 2 chooses L. Thus U is also a weakly dominated strategy. Formally,

Definition. A strategy $s_i \in S_i$ is called a weakly dominated strategy for player i in game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if there is another strategy $s'_i \in S_i$ such that

$$u_i(s'_i, s_{-i}) \geq u_i(s_i, s_{-i})$$

for all $s_{-i} \in S_{-i}$ with inequality holds for some s_{-i} . In this case we say that strategy s'_i weakly dominates s_i . A strategy is called a weakly dominant strategy for player i if it weakly dominates each of i 's other strategies.

In the game above, D is a weakly dominant strategy.

The second-price sealed bid auction (the vickrey auction): Each bidder submits a sealed bid. The bidder who bids the highest gets the object, and pays the price equal to the second-highest bid. Show that each bidder has a weakly dominant strategy: bidding his true valuation.

Proof. Let bidder i 's true valuation for the object be v_i , and his bid b_i . We show $b_i = v_i$ is a weakly dominant strategy for i . First consider any $b_i > v_i$. If b_i is not the winning bid, or if b_i wins and the second highest bid $b_{(2)}$, is less than v_i , then

bidding $b_i > v_i$ results in the same payoff to i as if he bids $b_i = v_i$. But if b_i wins and $b_{(2)} > v_i$, then i will lose the amount equal to $b_{(2)} - v_i$, while i could have avoided this loss by bidding v_i . Thus $b_i = v_i$ weakly dominates $b_i > v_i$. Next consider any $b_i < v_i$. If b_i is the winning bid, or if the highest bid is large than v_i , then $b_i < v_i$ result in the same payoff to i as if i bids $b_i = v_i$. But if the highest bid is higher than b_i but lower than v_i , then bidding $b_i < v_i$ yields zero payoff to i while bidding $b_i = v_i$ would have yielded positive payoff to i . Thus $b_i = v_i$ weakly dominates $b_i < v_i$. Thus $b_i = v_i$ is a weakly dominant strategy.

Unlike the case of strictly dominant strategy, it is less clear that a rational player should not choose a weakly dominated strategy. In the above example, if 1 is certain that 2 will choose L, then playing M is as good as playing D. So we need some additional consideration if a weakly dominated strategy is to be eliminated.

The logic behind the idea that a strictly dominated strategy should not be played can be extended in the following way. After a strictly dominated strategy is eliminated, we can consider what remains as a new game and again look for strictly dominated strategies in the new game. We can then eliminate any strictly dominated strategy in this game, and then start yet another new game, and so on. This process of achieving a prediction about what should not be played in game is called iterated deletion of strictly dominated strategies. Consider the next game, a modified version of the prisoner's dilemma game, called the DA's Brother.

| | | |
|------------|------------|---------|
| | Prisoner 2 | |
| | DC | C |
| DC | 0, -2 | -10, -1 |
| Prisoner 1 | | |
| C | -1, -10 | -5, -5 |

Now prisoner 1 has no dominant strategy: if 2 plays DC, 1's best response is DC; and if 2 plays C, 1's best response is C. But DC is a dominated strategy for 2, and if we eliminate this, then it becomes clear that 1 should play C. Thus the unique predicted outcome is still (C,C).

So far, we have considered only games with pure strategies. But the ideas can be easily extended to games allowing mixed strategies.

Definition. A strategy $\sigma_i \in \Delta(S_i)$ is strictly dominated for player i in game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ if there exists another strategy $\sigma'_i \in \Delta(S_i)$ such that for all $\sigma_{-i} \in \prod_{j \neq i} \Delta(S_j)$

$$u_i(\sigma'_i, \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i}).$$

Notice that since

$$u_i(\sigma'_i, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i}) = \sum_{s_{-i} \in S_{-i}} [\prod_{j \neq i} \sigma_j(s_j)] [u_i(\sigma'_i, s_{-i}) - u_i(\sigma_i, s_{-i})],$$

we have

$$u_i(\sigma'_i, \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i}) \text{ for all } \sigma_{-i} \text{ iff } [u_i(\sigma'_i, s_{-i}) - u_i(\sigma_i, s_{-i})] \text{ for all } s_{-i}.$$

Therefore to test whether a strategy σ_i is strictly dominated by σ'_i , we need only compare these two strategies payoffs to player i against all pure strategies of i's opponents.

It follows that a pure strategy s_i is strictly dominated in a game allowing mixed strategies ($\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$) if and only if there exists another strategy σ'_i such that for all $s_{-i} \in S_{-i}$,

$$u_i(\sigma'_i, s_{-i}) > u_i(s_i, s_{-i}).$$

It is possible that a strategy is not strictly dominated when only pure strategies are considered, but become strictly dominated when mixed strategies are allowed. Consider example:

| | | | |
|----------|---|----------|-------|
| | | Player 2 | |
| | | L | R |
| | U | 10, 1 | 0, 4 |
| Player 1 | M | 4, 2 | 4, 3 |
| | D | 0, 5 | 10, 2 |

Neither U nor D strictly dominates M. Thus M is not a strictly dominated strategy if only pure strategies are considered. But now consider a mixed strategy σ_1 that plays U or D each with probability 0.5. Then

$$\begin{aligned} u_1(\sigma_1, L) &= 5 > u_1(M, L) \text{ and} \\ u_1(\sigma_1, R) &= 5 > u_1(M, R). \end{aligned}$$

Thus M is strictly dominated by σ_1 .

Notice also that if a pure strategy s_i is strictly dominated, then any mixed strategy σ_i that plays s_i with positive probability is also strictly dominated. This is because if we let σ'_i strictly dominates s_i , and consider a strategy σ''_i that is the same as σ_i except that it plays σ'_i when s_i would have been played. Then i 's payoff from σ''_i will be strictly higher than from σ_i for any strategy profile of i 's opponents.

The iterated removal of strictly dominated strategies can also be performed when mixed strategies are allowed.

Rationalizable Strategies

Another way to think about how a game should be played is the idea of rationalizable strategies. The basic idea is the following: If a strategy for a player is the best response to some possible strategies of the player's opponents, then playing this strategy by this player can be rationalized. On the other hand, if a strategy for a player is never a best response, then it would not be rational to play such a strategy, and hence such a strategy should not be played. We now make this idea a little formal.

Definition 1 *In game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$, Strategy σ_i is a best response for player i to his opponents' strategies σ_{-i} if*

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i})$$

for all $\sigma'_i \in \Delta(S_i)$. Strategy σ_i is never a best response if there exists no σ_{-i} for which σ_i is a best response.

Clearly, a strategy that is never a best response should not be played. This can possibly eliminate some strategies in a game. But we can push this idea a little further. After these strategies that are never best responses are removed from a

game, we can look at the remaining game and eliminate strategies that are never best responses in the remaining game, and so on. Strategies that survive this process of iterated removal of strategies that are never a best response are called rationalizable strategies.

Notice that a strictly dominated strategy can never be a best response. This can help us eliminate strategies that are never a best response.

Consider example:

| | | | |
|---|------|------|--------|
| | l | m | r |
| U | 5, 7 | 5, 0 | 0, 1 |
| M | 7, 0 | 1, 7 | 0, 1 |
| D | 0, 0 | 0, 0 | 10, -1 |

First, r is never a best response since it is strictly dominated by a mixed strategy of playing l and m with equal probabilities. After r is removed, D can never be a best response since it is strictly dominated by M. After this, we can not eliminate any more strategy since l is the best response to U, m is the best response to M, U the br to m and M the br to l. Thus the set of rationalizable pure strategies for 1 is {U, M} and for 2 is {l, m}.

Two more things to note:

The set of a player's rationalizable strategies is no larger than the set of a player's strategies that survive the iterated removal of strictly dominated strategies. For a two-player game, these two sets are the same.

The set of rationalizable strategies does not depend on the order of removal of strategies that are never a best response.

Nash Equilibrium

The set of strategies remaining after iterated removal of strictly dominated strategies and the set of rationalizable strategies are often quite large, and thus often do

not yield precise predictions. In a sense, these two concepts are not very powerful. A strategy is justifiable according to these two concepts as long as there are some conjectures about the play of the game under which the strategy is optimal, whether or not the conjectured play of the game will be actually realized. The concept of Nash equilibrium requires that each player's strategy is a best response to the strategies actually played by his opponents. That is, a Nash equilibrium is a profile of strategies such that each player's strategy is optimal given the strategies of other players. Formally,

Definition 2 A strategy profile $s = (s_1, \dots, s_I)$ constitutes a Nash equilibrium of game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if for every $i = 1, 2, \dots, I$,

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$$

for all $s'_i \in S_i$.

Example:

| | l | m | r |
|---|------|------|------|
| U | 5, 3 | 0, 4 | 3, 5 |
| M | 4, 0 | 5, 5 | 4, 0 |
| D | 3, 5 | 0, 4 | 5, 3 |

The game has a unique pure-strategy Nash equilibrium: (M,m). Notice that if a strategy is a Nash equilibrium strategy, then it is not strictly dominated and it is rationalizable. But a strategy that is rationalizable need not be a Nash equilibrium strategy. In the above example, each strategy is a rationalizable strategy, and thus survives the iterated removal of strictly dominated strategies, but there is only one pair of strategies that is a Nash equilibrium.

We next consider an example where each player's strategy set is continuous.

The Cournot Competition: Suppose two firms, A, B, compete in quantities. That is, they choose q_A and q_B simultaneously. Suppose that market demand is given

by $Q = 100 - P$, and each firm's production cost is $C(q_i) = cq_i$. Find the Nash equilibrium of this game.

Firm A's payoff is

$$\begin{aligned}\pi_A(q_A, q_B) &= q_A(P - c) \\ &= q_A[100 - q_A - q_B - c]\end{aligned}$$

Given any q_B , the optimal q_A satisfies:

$$100 - 2q_A - q_B - c = 0.$$

Similarly, given any q_A , the optimal q_B satisfies:

$$100 - 2q_B - q_A - c = 0.$$

A Nash equilibrium is a pair of (q_A^*, q_B^*) that solves these two equations, and we have $(q_A^*, q_B^*) = (\frac{100-c}{3}, \frac{100-c}{3})$.

There can also be games that have more than one Nash equilibrium, as in the following example.

The Battle of the Sexes. A young couple need to decide what to do for an evening. They both want to spend the evening together, but the husband prefers to watch a boxing game, while the wife prefers to watch a movie:

| | Boxing | Movie |
|--------|--------|-------|
| Boxing | 1, 2 | 0, 0 |
| Movie | 0, 0 | 2, 1 |

There are two pure-strategy Nash equilibrium in the game: (Boxing, Boxing) and (Movie, Movie). It is not clear from the game itself what the predicted outcome will be. More generally, there is the issue of why we should use Nash equilibrium as a solution. I refer you to the discussion in the book on this issue.

Now re-consider the Matching Pennies game:

| | | |
|---|-------|-------|
| | H | T |
| H | -1, 1 | 1, -1 |
| T | 1, -1 | -1, 1 |

The game has no pure-strategy Nash equilibrium. But it does have a mixed strategy Nash equilibrium.

Definition 3 A mixed strategy profile $(\sigma_1, \sigma_2, \dots, \sigma_I)$ constitutes a Nash equilibrium in game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ if for all $i = 1, 2, \dots, I$,

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i})$$

for all $\sigma'_i \in \Delta(S_i)$.

In the Matching Pennies game, let's find a Mixed strategy Nash equilibrium. Suppose

$$\begin{aligned} \sigma_1(H) &= x, \quad \sigma_1(L) = 1 - x \\ \sigma_2(H) &= y, \quad \sigma_2(L) = 1 - y \end{aligned}$$

The equilibrium x and y can be determined as follows:

$$\begin{aligned} -x + (1 - x) &= x - (1 - x) \\ y - (1 - y) &= -y + (1 - y) \end{aligned}$$

We have $x = y = \frac{1}{2}$. The equilibrium probabilities are determined in such a way that makes each player indifferent between the two pure strategies. Now if a player is indifferent between the pure strategies, then it is optimal for him to randomize between them. Thus the resulting strategy profile is a Nash equilibrium. More generally,

Proposition 4 Let $S_i^+ \subset S_i$ denote the set of pure strategies that player i assigns positive probabilities in mixed strategy profile $\sigma = (\sigma_i, \sigma_{-i})$. σ is a Nash equilibrium if

and only if for all $i = 1, 2, \dots, I$,

$$(i) \quad u_i(s_i, \sigma_{-i}) = u_i(s'_i, \sigma_{-i}) \text{ for all } s_i, s'_i \in S_i^+$$

$$(ii) \quad u_i(s_i, \sigma_{-i}) \geq u_i(s'_i, \sigma_{-i}) \text{ for all } s_i \in S_i^+ \text{ and all } s'_i \notin S_i^+.$$

Proof. For necessity, note that if either (i) or (ii) is not satisfied, there will be some $s_i \in S_i^+$ and $s'_i \in S_i$ such that $u_i(s'_i, \sigma_{-i}) > u_i(s_i, \sigma_{-i})$. Consider a strategy σ'_i that differs from σ_i only in that s'_i will be played whenever s_i would have been played according to σ_i . Then $u_i(\sigma'_i, \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i})$. For sufficiency, suppose both (i) and (ii) are satisfied but σ is not a Nash equilibrium. Then there must be some σ'_i such that $u_i(\sigma'_i, \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i})$ for some i , which can be true only if there exists some $s'_i \in S_i$ such that $u_i(s'_i, \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i}) = u_i(s_i, \sigma_{-i})$ for all $s_i \in S_i^+$, contradiction. ■

Now going back to the Battle of Sexes game, which has two pure-strategy Nash equilibrium.

| | | |
|--------|--------|-------|
| | Boxing | Movie |
| Boxing | 1, 2 | 0, 0 |
| Movie | 0, 0 | 2, 1 |

The game also has a mixed strategy equilibrium, with the wife choose Boxing with probability $\frac{1}{3}$ and Movie with prob. $\frac{2}{3}$; and the husband choose Boxing with prob. $\frac{2}{3}$ and Movie with prob. $\frac{1}{3}$.

Notice that in a mixed strategy Nash equilibrium, a player is indifferent between two strategies that he plays with positive probabilities, so why does he bother to randomize anyway? The reason is an equilibrium consideration: if he does not randomize, then the other players will not be indifferent among the possible pure strategies and will not randomize.

Under fairly general conditions, a Nash equilibrium (possibly in mixed strategies) exists. Two existence results:

(1) Every Normal form game in which the set of each player's pure strategies is finite has a mixed strategy Nash equilibrium.

(2) Every Normal form game that allows only pure strategies has a Nash equilibrium if all $i = 1, 2, \dots, I$,

(i) S_i is a nonempty, convex, and compact subset of some Euclidean space R^M .

(ii) $u_i(s_1, \dots, s_I)$ is continuous in (s_1, \dots, s_I) and quasiconcave in s_i .

Games of incomplete Information and Bayesian Nash Equilibrium

In many situations, players do not have complete information about their opponents. A seller of the house or a car may not know the valuation of potential buyers, for example. These are situations of incomplete information. The approach to analyze these games is to convert the problem to one of imperfect information. In a game of imperfect information, although a player may not know the type of another player, every player in the game knows that it is a random realization from a known distribution. That is, a player's type is thought of being chosen by the nature. Thus the incomplete information about a player's type is changed to the imperfect information about the nature's choice. A player's type is only observed by the player himself. For instance, although a seller may not know the valuation of a buyer on a car, but the seller knows the buyer's valuation is a random draw from the uniform distribution on [\$10,000, \$15,000]. A game of imperfect information is also called a Bayesian game.

Consider a modified version of the DA's brother game. Suppose DA's brother has the same payoffs as before, but his opponent is one of the two types, type I and type II, each occurring with probability μ and $1 - \mu$. The type I prisoner has the same payoffs as before, but the type II prisoner hates to confess. The game is

| | | | | | |
|----|---------|---------|--|----|--------------------|
| | DC | C | | DC | C |
| DC | 0, -2 | -10, -1 | | DC | 0, -2 -10, -7 |
| C | -1, -10 | -5, -5 | | C | -1, -10 -5, -11 |

Player 1's possible strategies: C or DC.

Player 2's possible strategies:

C if type I, C if type II;

C if type I, DC if type II;

DC if type I, C if type II;

DC if type I, DC if type II.

Let's be a little formal about Bayesian games.

Think about a game where there are I players. Each player may have different types. Let Θ_i be player i 's type set, and $\theta_i \in \Theta_i$. The vector $(\theta_1, \dots, \theta_I)$ is a realization of I players' types. Let the joint distribution of $(\theta_1, \dots, \theta_I)$ be $F(\theta_1, \dots, \theta_I)$, which is assumed to be common knowledge. Let $\Theta = \Theta_1 \times \Theta_2 \times \dots \times \Theta_I$. Then Θ contains all possible vectors of $(\theta_1, \dots, \theta_I)$.

A player's strategy in a Bayesian game is a type-contingent plan that specifies her action for each possible realization of her types. Denote player i 's strategy by $s_i(\theta_i)$. Note that here θ_i should be considered as a type variable, and $s_i(\theta_i)$ defines i 's strategic choices as a function of her types. A profile of strategies in the game is therefore a vector of functions $(s_1(\theta_1), s_2(\theta_2), \dots, s_I(\theta_I))$. Sometimes it is also written as $(s_i(\theta_i), s_{-i}(\theta_{-i}))$, or $(s_i(\cdot), s_{-i}(\cdot))$. Let S_i be player i 's strategy set that contains all $s_i(\cdot)$, and $\{S_i\}$ the collection of all players' strategy sets.

The payoff for player i in a Bayesian game depends obviously on the strategies of all players in the game, as well as the realizations of types of each player. For any strategy profile and any realized $(\theta_1, \dots, \theta_I) \in \Theta$, we can calculate player i 's payoff (when her type is θ_i), and we denote this payoff as $u_i(s_i(\cdot), s_{-i}(\cdot), \theta_i)$, or $u_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i)$. It is important to note that this payoff function is defined over both strategy profiles

and vectors of type variables. That is, for each strategy profile and each realization of the types, there corresponds to a payoff value. The collection of all players' payoff functions is denoted by $\{u_i(\cdot)\}$.

A Bayesian game can then be denoted by $[I, \{S_i\}, \{u_i(\cdot)\}, \Theta, F(\cdot)]$.

Now, given any strategy profile $(s_1(\cdot), \dots, s_I(\cdot))$, we can calculate player i 's expected payoff under this strategy profile as

$$\begin{aligned}\tilde{u}_i(s_1(\cdot), \dots, s_I(\cdot)) &= E_{\theta}[u_i(s_1(\theta_1), \dots, s_I(\theta_I), \theta_i)] \\ &= \int \int \dots \int_{(\theta_1, \dots, \theta_I) \in \Theta} u_i(s_1(\theta_1), \dots, s_I(\theta_I), \theta_i) dF(\cdot).\end{aligned}$$

Then the Bayesian game is equivalent to a normal-form game $\Gamma_N = [I, \{S_i\}, \{\tilde{u}_i(\cdot)\}]$.

Definition. A (pure strategy) Bayesian Nash equilibrium for the Bayesian game $[I, \{S_i\}, \{u_i(\cdot)\}, \Theta, F(\cdot)]$ is a profile of decision rules $(s_1(\cdot), \dots, s_i(\cdot))$ that constitutes a Nash equilibrium of game $\Gamma_N = [I, \{S_i\}, \{\tilde{u}_i(\cdot)\}]$. That is, for every $i = 1, \dots, I$,

$$\tilde{u}_i(s_i, s_{-i}) \geq \tilde{u}_i(s'_i, s_{-i})$$

for all $s'_i \in S_i$.

One way to find a BNE in a Bayesian game is to use the definition of BNE. First, write down all strategies for each player. Second, calculate $\tilde{u}_i(s_i, s_{-i})$ for all s_i and s_{-i} and for all i . Finally, find (if there is any) $(s_1(\cdot), \dots, s_i(\cdot))$ that is a Nash equilibrium in game $\Gamma_N = [I, \{S_i\}, \{\tilde{u}_i(\cdot)\}]$. This method is especially useful when each player has only a few strategies and it is thus easy to construct the normal-form game. For instance, problem 8.E.1. in Problem set 3.

Under many situations, however, an alternative method of finding BNE may prove to be more convenient. The idea here is that for a strategy profile to be a BNE, it is necessary and sufficient that each player, for each of his possible types, is choosing an optimal response to the conditional distribution of opponents' strategies. Formally, we have:

Proposition. A strategy profile $(s_1(\cdot), s_2(\cdot), \dots, s_I(\cdot))$ is a Bayesian Nash equilibrium in Bayesian game $[I, \{S_i\}, \{u_i(\cdot), \Theta, F(\cdot)\}]$ if and only if, for all i and all $\bar{\theta}_i \in \Theta_i$ occurring with positive probability,

$$E_{\theta_{-i}}[u_i(s_i(\bar{\theta}_i), s_{-i}(\theta_{-i}), \bar{\theta}_i) \mid \bar{\theta}_i] \geq E_{\theta_{-i}}[u_i(s'_i(\bar{\theta}_i), s_{-i}(\theta_{-i}), \bar{\theta}_i) \mid \bar{\theta}_i]$$

for all $s'_i(\cdot) \in S_i$, where the expectation is taken over realizations of the other players' types conditional on player i 's type being $\bar{\theta}_i$.

Proof. For necessity, note that if the condition above does not hold for some i for some $\bar{\theta}_i \in \Theta_i$ that occurs with positive probability, then i can increase her expected payoff by changing her strategic choice in the event her type is $\bar{\theta}_i$, contradicting $(s_1(\cdot), s_2(\cdot), \dots, s_I(\cdot))$ being a BNE. For sufficiency, if the condition above holds for all $\bar{\theta}_i \in \Theta_i$ occurring with positive probability, then

$$\begin{aligned} \tilde{u}_i(s_i, s_{-i}) &= E_{\bar{\theta}_i \in \Theta_i} \{E_{\theta_{-i}}[u_i(s_i(\bar{\theta}_i), s_{-i}(\theta_{-i}), \bar{\theta}_i) \mid \bar{\theta}_i]\} \\ &\geq E_{\bar{\theta}_i \in \Theta_i} \{E_{\theta_{-i}}[u_i(s'_i(\bar{\theta}_i), s_{-i}(\theta_{-i}), \bar{\theta}_i) \mid \bar{\theta}_i]\} \\ &= \tilde{u}_i(s'_i, s_{-i}). \end{aligned}$$

for all s'_i and all i . Thus $(s_1(\cdot), s_2(\cdot), \dots, s_I(\cdot))$ is a BNE. ■

The DA's brother example: Player 2's optimal strategy is C if type I and DC if type II. For player 1, the expected payoffs from DC or C are: -10μ and $-5\mu - (1 - \mu)$.

Thus 1 should choose DC if

$$\mu < \frac{1}{6}$$

and choose C if $\mu > \frac{1}{6}$. The Bayesian Nash equ. is for 1 to choose DC if $\mu < \frac{1}{6}$ and C if $\mu > \frac{1}{6}$, and for 2 to choose C if type I and DC if type II.

The first-price sealed-bid auction: In a first-price sealed-bid auction for an object, each bidder hands in an envelop that indicates the amount he bids. The person who bids the highest wins the object and pays the amount he bids. Assume that each bidder's valuation for the object, v_i , is a random draw from UNIF $[0,1]$, if

there are only two bidders, find a symmetric Bayesian Nash equilibrium of bidding strategies.

Suppose both bidders use the same bidding function $b(v_i)$ that is monotonically increasing in v_i . $(b(v_1), b(v_2))$ is a BNE if, for $i \neq j$, given j 's strategy as $b(v_j)$, $b(v_i)$ is optimal for all v_i .

Player i 's expected payoff when his type is v_i but he bids as if his valuation is v will be

$$\begin{aligned} u_i(v, v_i) &= (v_i - b(v)) \Pr(b(v) > b(v_j)) \\ &= (v_i - b(v)) \Pr(v > v_j) \\ &= (v_i - b(v))v. \end{aligned}$$

Maximizing $u_i(v, v_i)$ with respect to v , we have:

$$\frac{\partial u_i}{\partial v} = v_i - b(v) - b'(v)v = 0.$$

Given $b(v_j)$, bidding $b(v_i)$ is optimal for any v_i if $u_i(v, v_i)$ is maximized when $v = v_i$, i.e., if

$$v_i - b(v_i) - b'(v_i)v_i = 0.$$

To solve this differential equation, guess $b(v_i) = \alpha v_i + \beta$, then

$$v_i - \alpha v_i - \beta - \alpha v_i = 0$$

which can hold for all v_i iff $\beta = 0$ and $\alpha = \frac{1}{2}$. Thus an Bayesian Nash equilibrium is for bidder i to bid according to the bidding function $b(v_i) = \frac{1}{2}v_i$.