Static Games of Incomplete Information

IE 5545: Decision Analysis
Prof. Krishnamurthy Iyer*

May 2, 2026

Until now we have focused on games of complete information, where the structure of the game and the rationality of players was common knowledge. If on the other hand this common knowledge assumption does not hold, then we obtain a game of incomplete information. This incomplete information can be of many forms: players may not know each others' payoff functions, or the actions they have, or whether the other players are rational.

We now describe a framework developed by Harsanyi to analyze games of incomplete information. The framework models one player's incomplete information about another player as the first player's *uncertainty* about the *type* of the second player. Adopting a Bayesian point of view, this uncertainty is captured via a probability distribution. The players' types and the probability distribution with which the types are chosen then convert the incomplete information game into an imperfect (but complete) information game, which can then be analyzed using tools we have already studied (Nash equilibrium, SPNE, etc).

Before presenting the framework, we illustrate it through an extended example.

1 Example: Spatial resource competition

Consider a setting involving players in a location, who each must decide whether to stay in that location or to move to a different one. Each location has a resource which is equally shared among everyone in that location. Thus, if a location has too many players, each will get a low reward due to resource competition, incentivizing them to move. At the same time, the decision to move incurs moving costs. Such spatial resource competition arises, for instance, among cab drivers in a city deciding where to position themselves to obtain rides.

We will consider a simple setting, with two players, each deciding between two actions $\mathsf{stay}(S)$ (in the current location) or $\mathsf{move}(M)$ (to a second location). Current location has a total resource of 4, whereas the second location has total resource of 6. Each player i incurs a cost of c_i for moving. Each player makes the decision to stay or move simultaneously. The payoff matrix is given by

	M	S
M	$3 - c_1, 3 - c_2$	$6 - c_1, 4$
S	$4, 6 - c_2$	$2, 2$

If the payoff matrix were common knowledge, then we have a static game with complete information. Assuming that $c_i \in [0, 4]$ (i.e., moving costs are not too high), we can immediately see that the game has three Nash equilibria:

1. pure NE $(\mathsf{stay}, \mathsf{move})$
2. pure NE $(\mathsf{move}, \mathsf{stay})$
3. mixed NE (σ_1, σ_2) where $\sigma_1 = [\frac{4-c_2}{5} \ \ \frac{1+c_2}{5}]^\top$ and $\sigma_2 = [\frac{4-c_1}{5} \ \ \frac{1+c_1}{5}]^\top$. (Here, the first component denotes the probability of choosing move.)

How reasonable are the above equilibria as explanation/prediction of outcomes in a realistic setting? Not very! On one hand, the pure NEs imply one player stays and another moves, regardless of the underlying costs. Thus, if we consider the pure NE $(\mathsf{stay}, \mathsf{move})$ then we would expect player 2 to move regardless of the underlying moving costs. We expect that moving costs play some role in players behavior, and hence in the outcome. On the other hand, if we look at the mixed Nash equilibrium, it says that how player 1 randomizes depends on player 2's moving cost (and vice versa). In particular, player 1's moving cost plays no role in her equilibrium strategy.

Looking back to our motivation, we expect the behavior of a player to depend on their own moving cost, but not on the other players' moving cost. The underlying reason for this is that while is reasonable to assume that each player knows their own moving cost, it is often not realistic that they know other player's moving cost. In other words, the costs (and the game payoffs) are not common knowledge. If we take this motivation seriously, we must analyze this setting under the assumption that players only know their own costs, i.e., in settings where the costs (and hence the game payoffs) are not common knowledge. This lands us in the realm of a static game with *incomplete information*. We will now introduce (through this example) the Bayesian framework to analyze such games.

Recall that in decision problems with incomplete information, we introduce a set of states to capture all possible realizations of the decision maker's incomplete information. Then, we consider the decision maker's beliefs over those states, and conclude that a rational Bayesian decision maker will choose an action that maximizes their expected payoffs (with respect to their beliefs). If before making a decision, the decision maker receives new information about the states, then she first updates her belief using Bayes' rule, and then chooses the action that maximizes her expected payoffs (now with respect to her posterior belief).

For static games of incomplete information, we take a similar path. In this context, the incomplete information is about player *types* (e.g., in our resource competition setting, is the other player a low-cost player or a high-cost player?). Hence, there is a set of types for each player, and a set of *type profiles*. Each player has a *common prior belief* over the type profiles, and after learning her own type, updates it to a posterior belief about others' types. Finally, given other players' strategies, each player chooses an action that maximizes her expected payoffs, given her own type and her beliefs (i.e., plays a *best response*). This strategy profile then constitutes a Nash equilibrium, which in this context is called a *Bayes-Nash equilibrium* (BNE). We now illustrate this framework in our spatial resource competition game.

Since the players' moving costs are not common knowledge, the *type profile* (equivalent to state in a decision problem) is appropriately defined as $\theta = (\theta_1, \theta_2) = (c_1, c_2)$. In other words, if θ is made common knowledge, we get back to the complete information game. Here, $\theta_i = c_i$ is referred to as the type of player i.

Suppose the players believe that the type profile takes values in the set $\Theta = \{(0,0), (0,2), (2,0), (2,2)\}$. In other words, each $\theta_i \in \{0, 2\} = \Theta_i$.

A *prior* belief μ is a distribution over Θ, with $\mu(\theta)$ being the belief that the type profile is θ. We assume this prior μ is shared by both players and this fact is common knowledge. (This is called the common prior assumption). For our setting, let us assume that the players' types are i.i.d., with

$\mathbf{P}(\theta_i = 2) = \beta$ and $\mathbf{P}(\theta_i = 0) = 1 - \beta$ for some $\beta \in [0, 1]$. We then have

$$\mu(0,0) = (1-\beta)^2, \qquad \mu(0,2) = (1-\beta)\beta,$$
$$\mu(2,0) = \beta(1-\beta), \qquad \mu(2,2) = \beta^2.$$

With these assumptions, we can now represent the static game between the players as a dynamic game, where Nature moves first and chooses a type profile $\theta \in \Theta$, and then reveals to each player i her type θ_i, after which the two players play the static game. If each player i chooses an action $a_i \in \{\mathsf{stay}, \mathsf{move}\}$, then the payoff of player i is given by $\pi_i(a_i, a_{-i}, \theta)$, where

$$\pi_1(\mathsf{stay}, \mathsf{stay}, \theta) = 2 \qquad \pi_1(\mathsf{stay}, \mathsf{move}, \theta) = 4$$
$$\pi_1(\mathsf{move}, \mathsf{stay}, \theta) = 6 - \theta_1 \qquad \pi_1(\mathsf{move}, \mathsf{move}, \theta) = 3 - \theta_1,$$

and player 2's payoff is defined analogously. The extensive form of this representation is shown in Figure 1.

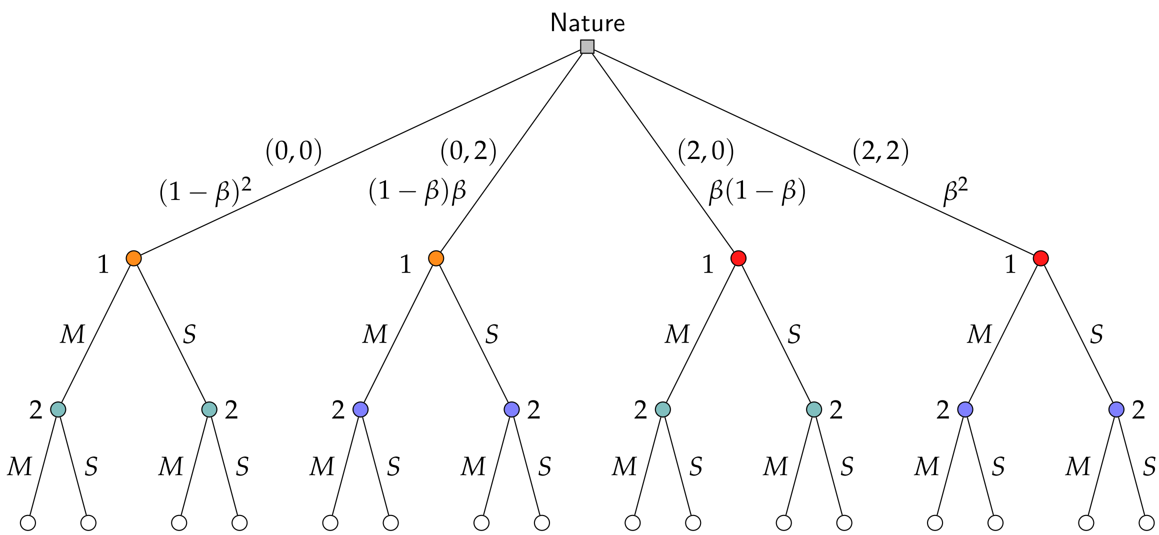

Figure 1: Extensive form for the resource competition game in the Bayesian framework. Nodes of the same color belong to the same information set.

In this extensive form, each player i has an information set corresponding to each possible type θ_i that she can be. Thus, a pure strategy s_i for a player i specifies which action the player should take for each possible type. In other words, a pure strategy for a player i is a function that maps each type $\theta_i \in \{0, 2\}$ to an action $\{\mathsf{move}, \mathsf{stay}\}$. For example, one pure strategy for player 1 involves staying when the costs are high and moving when costs are low, and is given by

$$s_1(0) = \mathsf{move}, \qquad s_1(2) = \mathsf{stay},$$

Similarly, another pure strategy which involves moving regardless of the type is given by

$$s_1(0) = \mathsf{move}, \qquad s_1(2) = \mathsf{move}.$$

1.1 A Bayes-Nash equilibrium

Suppose we want know if the strategy profile (s_1, s_2) is a (Bayes)-Nash equilibrium, where each s_i is given by

$$s_i(0) = \mathsf{move} \qquad s_i(2) = \mathsf{stay}.$$

To this, we need to check whether players are playing best responses to each other. Due to symmetry, it is enough to fix s_2 and verify whether player 1's strategy s_1 is a best response. To do this, we need to look at each information set (i.e., each possible value of θ_1) and ask whether the action suggested by s_1 achieves highest expected payoff among all actions.

Suppose $\theta_1 = 0$. Then, player 1's expected payoff from playing move is given by

$$\begin{aligned}
&\Phi_1(\mathsf{move}|\theta_1 = 0, s_2) \\
&\quad = (3 - \theta_1) \cdot \mathbf{P}(\text{player 2 chooses } \mathsf{move}) + (6 - \theta_1) \cdot \mathbf{P}(\text{player 2 chooses } \mathsf{stay}) \\
&\quad = 3 \cdot \mathbf{P}(s_2(\theta_2) = \mathsf{move}) + 6 \cdot \mathbf{P}(s_2(\theta_2) = \mathsf{stay}) \\
&\quad = 3 \cdot \mathbf{P}(\theta_2 = 0) + 6 \cdot \mathbf{P}(\theta_2 = 2) \\
&\quad = 3(1 - \beta) + 6\beta \\
&\quad = 3 + 3\beta.
\end{aligned}$$

Here, the third equality follows because $s_2(\theta_2) = \mathsf{move}$ if and only if $\theta_2 = 0$. In the penultimate equality, we use the fact that $\mathbf{P}(\theta_2 = 2) = \beta$ according to the common prior μ.

Similarly, player 1's expected payoff from playing stay is given by

$$\begin{aligned}
&\Phi_1(\mathsf{stay}|\theta_1 = 0, s_2) \\
&\quad = 4 \cdot \mathbf{P}(\text{player 2 chooses } \mathsf{move}) + 2 \cdot \mathbf{P}(\text{player 2 chooses } \mathsf{stay}) \\
&\quad = 4 \cdot \mathbf{P}(s_2(\theta_2) = \mathsf{move}) + 2 \cdot \mathbf{P}(s_2(\theta_2) = \mathsf{stay}) \\
&\quad = 4 \cdot \mathbf{P}(\theta_2 = 0) + 2 \cdot \mathbf{P}(\theta_2 = 2) \\
&\quad = 4(1 - \beta) + 2\beta \\
&\quad = 4 - 2\beta.
\end{aligned}$$

Since $s_1(0) = \mathsf{move}$, for s_1 to be a best response, we must have the expected payoff under move to be greater than (or equal to) that under stay, i.e., $\Phi_1(\mathsf{move}|\theta_1 = 0, s_2) \geq \Phi_1(\mathsf{stay}|\theta_1 = 0, s_2)$. This implies that $3 + 3\beta \geq 4 - 2\beta$, and hence we must have $\beta \geq \frac{1}{5}$.

Similarly, for $\theta_1 = 2$, we obtain player 1's expected payoffs to be

$$\begin{aligned}
\Phi_1(\mathsf{move}|\theta_1 = 2, s_2) &= (3 - 2) \cdot \mathbf{P}(s_2(\theta_2) = \mathsf{move}) + (6 - 2) \cdot \mathbf{P}(s_2(\theta_2) = \mathsf{stay}) \\
&= 1 \cdot \mathbf{P}(\theta_2 = 0) + 4 \cdot \mathbf{P}(\theta_2 = 2) \\
&= 1 - \beta + 4\beta \\
&= 1 + 3\beta \\
\Phi_1(\mathsf{stay}|\theta_1 = 2, s_2) &= 4 - 2\beta.
\end{aligned}$$

Since $s_1(2) = \mathsf{stay}$, we need $\Phi_1(\mathsf{stay}|\theta_1 = 2, s_2) \geq \Phi_1(\mathsf{move}|\theta_1 = 2, s_2)$, implying $\beta \leq \frac{3}{5}$.

Thus, if $\beta \in [\frac{1}{5}, \frac{3}{5}]$, the strategy s_1 is a best response for player 1 against s_2. Similarly, by symmetry, player 2 plays a best response by playing s_2 against s_1. Since both players are playing a best response we conclude that if $\beta \in [\frac{1}{5}, \frac{3}{5}]$ is a Bayes-Nash equilibrium.

The strategies s_1 and s_2 are shown in Figure 2, along with the resulting equilibrium paths for each realization of the type profile.

Now suppose $\beta = \frac{1}{2}$, so that indeed (s_1, s_2) is a pure Bayes-Nash equilibrium. What is the expected payoff of each player under this equilibrium? Knowing the strategies, this can be computed easily by following each equilibrium path in Figure 2, multiplying the resulting payoff with the probability of that path, and summing the resulting values over all paths. Denoting

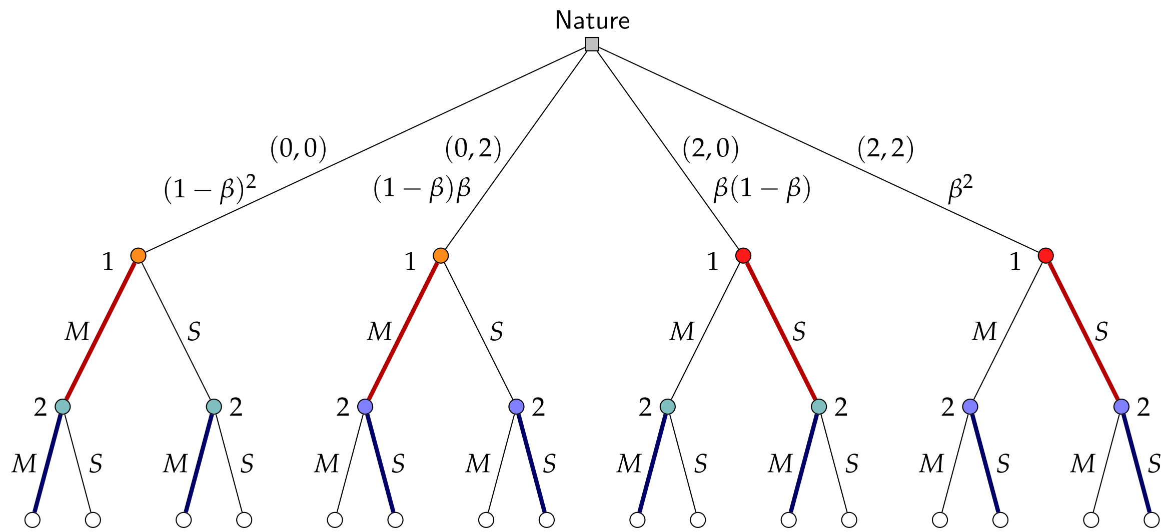

Figure 2: The strategies s_1 (red) and s_2 (blue) shown in bold. Traversing the bold lines yields an equilibrium path for each realization of the type profile. For instance, if the realized type profile is $(0,0)$, then the outcome is (M,M) and the resulting payoffs are $(3,3)$.

$\Pi_i(s_1, s_2)$ as player i's payoff under (s_1, s_2), we obtain

$$\begin{aligned}\Pi_1(s_1, s_2) &= \mu(0,0)\cdot 3 + \mu(0,2)\cdot 4 + \mu(2,0)\cdot 6 + \mu(2,2)\cdot 2\\ &= (1-\beta)^2\cdot 3 + (1-\beta)\beta\cdot 4 + \beta(1-\beta)\cdot 6 + \beta^2\cdot 2\\ &= \frac{15}{4}.\end{aligned}$$

(Similarly, we obtain $\Pi_2(s_1, s_2) = \frac{15}{4}$.) Here, we have taken expectation over both (θ_1, θ_2), so this is the expected payoff of the player 1 *before* Nature reveals θ_1 to player 1. For this reason, we term this as player 1's *ex ante* expected payoffs. (Ex ante is Latin for "before"). Note that *ex ante* expected payoff is a function of the strategy profile, but not of the types.

We can similarly ask for player 1's expected payoff for each action *after* her type is realized, but before she chooses her action (and before she knows player 2's type or her actions). This is nothing but $\Phi_1(a_1|\theta_1, s_2)$ that we computed while checking if (s_1, s_2) is a BNE. Since this expectation is taken at an intermediate stage where player 1 knows her type but not player 2's type, it is called player 1's *interim* expected payoff. A player's interim expected payoff is a function of her action, her type and other players' strategies.

Finally, the players' payoff at the end of the game, when actions and types of all players is known is just $\pi_i(a_i, a_{-i}, \theta_i, \theta_{-i})$. In analogy to others, this payoff is also called *ex post* payoff (Latin for "after"). Note that a player's ex post payoff is a function of the actions and types of all players.

2 General Bayesian framework

We now formally present the framework illustrated in the previous section for a general static game of incomplete information.

In the Bayesian framework, a static game of incomplete information consists of the following data:

1. A set of players $\{1, 2, \ldots, N\}$.
2. For each player i, a set Θ_i consisting of each possible *type* of player i.

   We let $\theta_i \in \Theta_i$ denote the type of player i, and $\theta = (\theta_1, \ldots, \theta_N)$ denote the *type profile*. Let Θ denote the set of all possible type profiles. Note that $\Theta = \Theta_1 \times \Theta_2 \times \cdots \times \Theta_N$.

   Similarly, we let $\theta_{-i} = (\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_N)$ denote the types of every player other than player i, and define Θ_{-i} analogously.
3. A probability distribution $\mu \in \Delta(\Theta)$ over the set of all type profiles Θ.

   This distribution μ describes how each player *believes* the type profile θ is distributed. Underlying this statement is the assumption that each player has the *same* belief μ. We further assume that this fact is common knowledge among the players. Together, these two assumptions constitute the *common prior* assumption.

   Note that **P** can be any distribution; in particular, the types of players i and j need not be independent.
4. For each player i, a set of actions $\mathcal{A}_i$.

   We define $a = (a_1, \ldots, a_N)$ as an action profile, where $a_i \in \mathcal{A}_i$ for each i. Similarly, we define a_{-i}. We assume each player i simultaneously chooses an action $a_i \in \mathcal{A}_i$.
5. A payoff function π_i for each player i. Here $\pi_i(a_i, a_{-i}, \theta_i, \theta_{-i})$ denotes the payoff of player i when she plays a_i, has type θ_i, all other players play a_{-i} and have types θ_{-i}.

Finally, we assume the above description is common knowledge among the players. In other words, each player i knows who the players are, what each player j's set of possible types Θ_j is, how the types of each player is chosen, what actions the players have, and what payoffs they receive as a function of their actions and their types. Moreover, each player knows that every other player knows this, each player knows that every other player knows that every other player knows this, and so on.

2.1 How the game unfolds

The sequence of events in a static game of incomplete information is as follows:

1. Nature first randomly chooses a type profile $\theta = (\theta_1, \ldots, \theta_N)$ according to the probability distribution μ.
2. Nature then reveals to each player i her type θ_i. The types θ_{-i} of all other players are kept hidden from player i.
3. All players i, knowing only their own types θ_i, simultaneously choose actions $a_i \in \mathcal{A}_i$.
4. Each player i receives a payoff $\pi_i(a_i, a_{-i}, \theta_i, \theta_{-i})$.

2.2 Pure strategies

As we have defined, each player i has an information set corresponding to each value of $\theta_i \in \Theta_i$ that Nature reveals to her. Hence, a pure strategy $s_i(\cdot)$ for player i must specify an action $s_i(\theta_i) \in \mathcal{A}_i$ for each possible realization of her type θ_i. In other words, we can specify a pure strategy $s_i : \Theta_i \to \mathcal{A}_i$

for player i by specifying which action $s_i(\theta_i)$ she will take when her type is θ_i, for different values of θ_i.

Note that an action a_i is an element of the set $\mathcal{A}_i$, whereas a strategy $s_i(\cdot)$ is a function that takes values in Θ_i and outputs an action in $\mathcal{A}_i$. In other words, for each θ_i, $s_i(\theta_i)$ is an action, whereas s_i itself is a function $s_i : \Theta_i \to \mathcal{A}_i$. In order to make the distinction clearer, we will as far as possible denote an action by the letter a and a strategy by the letter s, with appropriate subscripts for the player.

2.3 Expected payoffs

There are multiple points during the game, where a player might evaluate her payoffs corresponding to different strategies or different actions available to her.

1. Before playing the game and before Nature reveals the type to a player, she may want to know her expected payoff for playing different strategies. This stage is referred to as the *ex ante* stage.

2. Alternatively, a player, after knowing her type, may want to evaluate her payoffs for choosing different actions. This stage is referred to as the *interim* stage, since it is evaluated during the game — after a player knows her type, but before she chooses her action.

3. For the sake of completeness, $\pi_i(a_i, a_{-i}, \theta_i, \theta_{-i})$ is called *ex post* payoff – this is the expected payoff of player i after she knows her type and those of other players, and everyone's actions.

In order to evaluate a strategy at the *ex ante* stage, or an action at the *interim* stage, we will need to consider a players' expectation of her (*ex post*) payoff at that stage. We define these concepts below.

1. Suppose each player i adopts a pure strategy $s_i : \Theta_i \to \mathcal{A}_i$. Then, the expected payoff of player i is given by

$$\Pi_i(s_1, \ldots, s_N) = \mathbf{E}\left[\pi_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i})\right] = \sum_{\theta \in \Theta} \mu(\theta) \cdot \pi_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}).$$

   Note that since the player does not know her type as well as others, the expectation here is over the entire type profile $\theta = (\theta_i, \theta_{-i})$. This is the expected payoff of the player i **before** she knows what her type is, based only on her and others' strategies. This is also referred to as *ex ante* payoff.

2. Next, suppose each player $j \neq i$ follows a pure strategy $s_j : \Theta_j \to \mathcal{A}_j$. Suppose player i has type θ_i, and chooses an action a_i. Then, her expected payoff after she knows her type is given by

$$\Phi_i(a_i|\theta_i, s_{-i}) = \mathbf{E}\left[\pi_i(a_i, s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}|\, \theta_i\right] = \sum_{\theta_{-i} \in \Theta_{-i}} \mu(\theta_{-i}|\, \theta_i) \cdot \pi_i(a_i, s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}),$$

   where $\mu(\theta_{-i}|\theta_i)$ denotes the *conditional probability* that other players' type profile is θ_{-i}, given player i's type is θ_i. This conditional probability is given by:

$$\mu(\theta_{-i}|\theta_i) = \frac{\mu(\theta_i, \theta_{-i})}{\sum_{\widehat{\theta}_{-i} \in \Theta_{-i}} \mu(\theta_i, \widehat{\theta}_{-i})},$$

assuming the denominator is positive.

Now, the preceding expression denotes player i's expected payoff **after** she knows her type and decides to choose an action a_i, but **before** she knows what other players' types and actions are. Consequently, the expectation is only over θ_{-i}. This expected payoff is also known as *interim* expected payoff.

Note that the interim expected payoff $\Phi_i(a_i|\theta_i, s_{-i})$ of a player i depends only on her type θ_i, the action a_i she chooses and the strategies of all other players s_{-i}. When the other players' strategies are fixed, we sometimes drop it from the notation and write the interim payoffs as $\Phi_i(a_i|\theta_i)$, to keep the notation light.

It is a good exercise to write down the relation between the *ex ante* payoff Π_i and the interim payoff Φ_i. (Exercise: write down this relation.)

2.4 Bayes-Nash equilibrium

Given the ex ante payoffs Π_i of each player i for each strategy profile, we obtain a game in strategic form. For this strategic form, we can write down the definition of a Nash equilibrium as before. Since we have adopted a Bayesian framework, Nash equilibrium in this setting is referred to as a *Bayes-Nash equilibrium*, or Bayesian Nash equilibrium, or simply a Bayesian equilibrium. We will use the terminology Bayes-Nash equilibrium. We obtain the following definition:

Definition 1 (Bayes-Nash equilibrium (BNE))**.** *A strategy profile $(s_1, s_2, \ldots, s_N)$ is a (pure) Bayes-Nash equilibrium, if for each player i, we have*

$$\Pi_i(s_i, s_{-i}) \geq \Pi_i(s'_i, s_{-i}), \quad \textit{for all strategies } s'_i : \Theta_i \to \mathcal{A}_i.$$

Since the former definition involves maximizing over pure strategies, which are functions, we look at an alternative but equivalent definition, involving interim payoffs.

Suppose each player $j \neq i$ follows the strategy $s_j : \Theta_j \to \mathcal{A}_j$. Now, for s_i to be a best response for player i, it must the case that the optimal action for player i to play is $s_i(\theta_i)$ if she observes her type is θ_i. In other words, it must be true that

$$\Phi_i(s_i(\theta_i)|\theta_i, s_{-i}) \geq \Phi_i(a_i|\theta_i, s_{-i}), \quad \text{for all } a_i \in \mathcal{A}_i \text{ and for all } \theta_i \in \Theta_i.$$

This leads to this equivalent second definition:

Definition 2 (Bayes-Nash equilibrium (BNE))**.** *A strategy profile $(s_1, s_2, \ldots, s_N)$ is a (pure) Bayes-Nash equilibrium, if for each player i, we have*

$$\Phi_i(s_i(\theta_i)|\theta_i, s_{-i}) \geq \Phi_i(a_i|\theta_i, s_{-i}), \quad \textit{for all } a_i \in \mathcal{A}_i \textit{ and for all } \theta_i \in \Theta_i.$$

Often this second definition is analytically (and computationally) simpler. (This second definition is what we adopted in the example in the preceding section.)

2.4.1 Exercise: Continuous types

Consider the spatial resource competition game, but now assume that, under the common prior μ, the type θ_i of each player i is distributed independently and identically as $\mathsf{Unif}[0, 2]$. Find a *symmetric* Bayes-Nash equilibrium where the players use strategies (s_1, s_2)

$$s_1(x) = s_2(x) = \begin{cases} \mathsf{move} & \text{if } x \leq c_0 \\ \mathsf{stay} & \text{if } x > c_0 \end{cases}$$

for some $c_0 \in [0, 2]$.

Hint: Fix s_2 as above. Compute player 1's interim expected payoff $\Phi_1(a_1|\theta_1, s_2)$ for each value of $a_1 \in \{\mathsf{stay}, \mathsf{move}\}$ and $\theta_1 \in [0, 2]$. You are looking for a value of c_0 such that if $\theta_1 \leq c_0$, then $\Phi_1(\mathsf{move}|\theta_1, s_2) \geq \Phi_1(\mathsf{stay}|\theta_1, s_2)$, and if $\theta_1 > c_0$, then $\Phi_1(\mathsf{move}|\theta_1, s_2) \leq \Phi_1(\mathsf{stay}|\theta_1, s_2)$.

