

Bayesian Games

CSCI 1440/2440

2024-01-31

We describe incomplete-information, or Bayesian, normal-form games (formally; no examples), and corresponding equilibrium concepts.

1 A Bayesian Model of Interaction

A **Bayesian**, or **incomplete-information**, game is a generalization of a complete-information game. Recall that in a complete-information game, everything relevant to the play of the game is assumed to be common knowledge. In a Bayesian game, many things about the game are again common knowledge, but the players may have additional private information. This private information is captured by the notion of an **epistemic type**, which describes a player's private knowledge: i.e., all that is not common knowledge. The Bayesian-game formalism makes two simplifying assumptions:

- Any information that is not common knowledge pertains only to utilities. In particular, in all realizations of a Bayesian game, the number of players and each one's action set across realizations are identical.¹
- Players maintain beliefs about the game (i.e., about utilities) in the form of a joint probability distribution over all players' types. Prior to receiving any private information, this probability distribution is common knowledge. As such, it is called a **common prior**.²

After receiving private information, players condition on this newfound knowledge to update their beliefs. As a consequence of the common prior assumption, any differences in beliefs can be attributed entirely to differences in information.

As in complete-information games, rational players are assumed to maximize their expected utilities. Further, they are assumed to update their beliefs via Bayes' rule when they learn new information.

Thus, in a Bayesian game, in addition to players, actions, and utilities, there is a type space $T = \prod_{i \in [n]} T_i$, where T_i is the type space of player i . There is also a common prior F , which is a probability distribution over type profiles. Utility functions then depend not only on actions, but on types as well: $u_i : A \times T \rightarrow \mathbb{R}$.

Again, as in complete-information games, it is assumed that all of the above is common knowledge among the players. Further, it is assumed that each player learns her own type: i.e., receives the

¹ This is not a strong assumption. If this were not the case, e.g., if one player were unsure as to whether another player had one or two actions, a dummy action (e.g., a dominated strategy) could be added to a model of the game in the first case, so that the other player always has two actions.

² This is a strong assumption.

relevant private information. An agent's strategy, then, becomes a function of its type: i.e., for all players i , $s_i : T_i \rightarrow A_i$. And a mixed strategy, as usual, is a probability distribution over (pure) strategies.³

³ Mixed strategies in Bayesian games are complicated objects: they are probability distributions over functions!

2 Phases of a Bayesian Game

We can divide a Bayesian game into three phases:

- In the **ex-ante** phase, no player knows what her own type is, or the types of any other player. When we reason about a strategy in this phase, we use the following expected utility function:

$$\mathbb{E}_{\mathbf{t} \sim F} [u_i(s_i, \mathbf{s}_{-i}; t_i, \mathbf{t}_{-i})].$$

- In the **interim** phase, each player knows her own type, but not the types of any other player. When we reason about a strategy in this phase, we use the following expected utility function:

$$\mathbb{E}_{\mathbf{t}_{-i} \sim F_{\mathbf{t}_{-i}|t_i}} [u_i(s_i, \mathbf{s}_{-i}; t_i, \mathbf{t}_{-i})].$$

The notation $F_{\mathbf{t}_{-i}|t_i}$ signifies the joint distribution over type profiles conditioned on i 's private information.

- In the **ex-post** phase, each player knows her own type, and the types of every other player. When we reason about a strategy in this phase, we use the following utility function:

$$u_i(s_i, \mathbf{s}_{-i}; t_i, \mathbf{t}_{-i}).$$

We summarize the phases in Figure 1:

Phase	Knows t_i ?	Knows \mathbf{t}_{-i} ?	Relevant Utility
Ex-ante	No	No	$\mathbb{E}_{\mathbf{t} \sim F} [u_i(s_i, \mathbf{s}_{-i}, t_i, \mathbf{t}_{-i})]$
Interim	Yes	No	$\mathbb{E}_{\mathbf{t}_{-i} \sim F_{\mathbf{t}_{-i} t_i}} [u_i(s_i, \mathbf{s}_{-i}, t_i, \mathbf{t}_{-i})]$
Ex-post	Yes	Yes	$u_i(s_i, \mathbf{s}_{-i}, t_i, \mathbf{t}_{-i})$

Figure 1: A summary of the phases of a Bayesian Game.

3 Bayesian Equilibria in Bayesian Games

Corresponding to the three phases of a Bayesian game, there are three notions of equilibrium in a Bayesian game.

A strategy profile $\mathbf{s} = (s_i, \mathbf{s}_{-i}) \in S$ is an **ex-ante Bayes-Nash equilibrium** if no player can increase her ex-ante expected utility by unilaterally changing her strategy:

$$\mathbb{E}_{\mathbf{t} \sim F} [u_i(s_i(t_i), \mathbf{s}_{-i}(t_{-i}); t_i, \mathbf{t}_{-i})] \geq \mathbb{E}_{\mathbf{t} \sim F} [u_i(s'_i(t_i), \mathbf{s}_{-i}(t_{-i}); t_i, \mathbf{t}_{-i})], \quad \forall i \in [n], \forall s'_i \in S_i.$$

A strategy profile $\mathbf{s} = (s_i, \mathbf{s}_{-i}) \in S$ is an **interim Bayes-Nash equilibrium** if no player can increase her interim expected utility by unilaterally changing her strategy: $\forall i \in [n], \forall t_i \in T_i, \forall s'_i \in S_i$,

$$\mathbb{E}_{\mathbf{t}_{-i} \sim F_{\mathbf{t}_{-i}|t_i}} [u_i(s_i(t_i), \mathbf{s}_{-i}(t_{-i}); t_i, \mathbf{t}_{-i})] \geq \mathbb{E}_{\mathbf{t}_{-i} \sim F_{\mathbf{t}_{-i}|t_i}} [u_i(s'_i(t_i), \mathbf{s}_{-i}(t_{-i}); t_i, \mathbf{t}_{-i})].$$

Interestingly, ex-ante and interim Bayes-Nash equilibria turn out to be equivalent, in which case they are both referred to merely as **Bayes-Nash equilibria** (BNE). Moreover, as Bayes-Nash equilibria are Nash equilibria (imagine exploding a Bayesian game into a normal-form game; example forthcoming), Nash's theorem guarantees their existence, assuming T is finite.

Proof of Claim: Interim equilibrium is a stronger notion than ex-ante equilibrium: by simply taking expectations over t_i , we recover the ex-ante equilibrium condition from the interim one. Thus, it remains only to establish the other direction, that ex-ante equilibria are also interim equilibria. Recall that strategies are functions from types to actions. It suffices to argue that any strategy that is a best response in expectation across types necessarily generates a best response for each specific type. Suppose otherwise: that there is a strategy that is a best response in expectation across all types but does *not* generate a best response for some specific type, say t_i . Then this strategy is not in fact a best response in expectation across all types, since it could generate a better response for t_i .

4 Ex-post Nash Equilibria in Bayesian Games

A strategy profile $\mathbf{s} = (s_i, \mathbf{s}_{-i}) \in S$ is an **ex-post Nash equilibrium** (EPNE) if no player can increase her *ex-post* expected utility by unilaterally changing her strategy:

$$u_i(s_i(t_i), \mathbf{s}_{-i}(t_{-i}); t_i, \mathbf{t}_{-i}) \geq u_i(s'_i(t_i), \mathbf{s}_{-i}(t_{-i}); t_i, \mathbf{t}_{-i}), \quad \forall i \in [n], \forall s'_i \in S_i, \forall \mathbf{t} \in T.$$

EPNE is an equilibrium concept. As such, it assumes all players are best responding to one another: i.e., maximizing their expected utility, given the other players' strategies. In a sense, however, it is a worst-case concept, because it does not assume the other players are Bayesian. Rather, each player's strategy, which is conditioned on her own type, must be a best response to the others' strategies, which are also conditioned on their types, *regardless of the other players' types*.

Taking this worst-case reasoning one step further, it is also possible to define DSE in incomplete-information games, by dropping both assumptions—that players are *expected* utility maximizers *and* that they are Bayesian! As in complete-information games, DSE need not exist in incomplete-information games (a strict generalization).

Formally, a strategy s_i for player $i \in [n]$ in an incomplete-information game is **dominant** if it is (weakly) optimal, regardless of the other players' actions *and types*: i.e.,

$$u_i(s_i(t_i), \mathbf{a}_{-i}; t_i, \mathbf{t}_{-i}) \geq u_i(s'_i(t_i), \mathbf{a}_{-i}; t_i, \mathbf{t}_{-i}), \quad \forall \mathbf{t} \in T, \forall s'_i \in S_i, \forall \mathbf{a}_{-i} \in A_{-i}$$

A strategy profile $\mathbf{s} \in S$ is a **dominant strategy equilibrium (DSE)** if all players play dominant strategies.

Like DSE, EPNE need not exist. To make this point, we present three examples, two games with EPNE, and a final game that combines the prior two, but does not have an EPNE.⁴

⁴ This counterexample was borrowed from these lecture notes.

Example 4.1 (Example of an EPNE). Consider a two-player Bayesian game with type space $T_1 = \{T\}$ and $T_2 = \{L, R\}$ and action spaces $A_1 = A_2 = \{C, D\}$. The payoffs of this game are described by the following two matrices, or subgames, each one corresponding to a possible type profile: i.e., TL and TR .

		L			R	
		C	D	C	D	
T	C	2,2	0,0	C	2,1	0,0
	D	3,0	1,1*	D	3,0	1,2*

Observe that these two subgames both have a pure-strategy Nash equilibrium profile, indicated by a *. Moreover, these equilibria are dominant (and hence, pure) strategy equilibria in their respective subgames, so they are the unique equilibria (mixed or pure) of these subgames. Consequently, in any ex-post equilibrium, players must play one of these two equilibrium profiles.

Player 2's actions in these equilibria cannot depend on player 1's type, as player 1 has only 1 possible type. So the interesting case to consider is player 1's actions; specifically, whether they vary with player 2's type. As they do not—player 1 plays D regardless of player 2's type—this game affords a unique EPNE, namely D, D .

Example 4.2 (Another example of an EPNE). Consider a two-player Bayesian game with type space $T_1 = \{B\}$ and $T_2 = \{L, R\}$ and action spaces $A_1 = A_2 = \{C, D\}$. The payoffs of this game are described by the following two matrices, or subgames, each one corresponding to a possible type profile: i.e., BL and BR .

		L			R	
		C	D	C	D	
B	C	1,2*	3,0	C	1,1*	3,0
	D	0,0	2,1	D	0,0	2,2

As in the previous example, these two subgames both have a pure-strategy Nash equilibrium profile, indicated by a *. Moreover, these equilibria are dominant (and hence, pure) strategy equilibria in their

respective subgames, so they are the unique equilibria (mixed or pure) of these subgames. Consequently, in any ex-post equilibrium, players must play one of these two equilibrium profiles.

Player 2's actions in these equilibria cannot depend on player 1's type, as player 1 has only 1 possible type. So the interesting case to consider is player 1's actions; specifically, whether they vary with player 2's type. As they do not—player 1 plays *C* regardless of player 2's type—this game affords a unique EPNE, namely *C, C*.

Example 4.3 (Counterexample to the existence of EPNE). Consider a two-player Bayesian game with type spaces $T_1 = \{T, B\}$ and $T_2 = \{L, R\}$ and action spaces $A_1 = A_2 = \{C, D\}$. The payoffs of this game are described by the following four matrices, or subgames, each one corresponding to a possible type profile: i.e., *TL, TR, BL, or BR*.

		<i>L</i>		<i>R</i>	
		<i>C</i>	<i>D</i>	<i>C</i>	<i>D</i>
<i>T</i>	<i>C</i>	2,2	0,0	<i>C</i>	2,1
	<i>D</i>	3,0	1,1*	<i>D</i>	3,0
		<i>C</i>	<i>D</i>	<i>C</i>	<i>D</i>
<i>B</i>	<i>C</i>	1,2*	3,0	<i>C</i>	1,1*
	<i>D</i>	0,0	2,1	<i>D</i>	0,0

As above, each of these four subgames has a pure-strategy Nash equilibrium profile, indicated by a *. Moreover, these equilibria are dominant (and hence, pure) strategy equilibria in their respective subgames, so they are the unique equilibria (mixed or pure) of these subgames. Consequently, in any ex-post equilibrium, players must play one of these four equilibrium profiles.

The EPNE constraints do not pose a problem for player 1. When player 1 is of type *T* (as in Example 4.1), player 1 plays *D* in both equilibria, regardless of whether player 2 is of type *L* or *R*. Likewise, when player 1 is of type *B* (as in Example 4.2), player 1 plays *C* in both equilibria, regardless of whether player 2 is of type *L* or *R*.

But player 2 cannot satisfy the EPNE constraints. Regardless of player 2's type, she would have to play *D* when player 1 is of type *T*, and *C* when player 1 is of type *B*. As players cannot condition their play on *another* player's type, no EPNE exist in this game.