14.126 GAME THEORY

MIHAI MANEA

Department of Economics, MIT,

1. Normal-Form Games

A **normal** (or **strategic**) **form game** is a triplet (N, S, U) with the following properties

- $N = \{1, 2, \ldots, n\}$ is a finite set of **players**
- S_i is the set of **pure strategies** of player i; $S = S_1 \times \ldots \times S_n$
- $u_i : S \to \mathbb{R}$ is the **payoff function** of player i; $u = (u_1, \ldots, u_n)$.

Player $i \in N$ receives payoff $u_i(s)$ when $s \in S$ is played. The game is **finite** if S is finite.

The structure of the game is **common knoweldge**: all players know (N, S, U), and know that their opponents know it, and know that their opponents know that they know, and so on.

1.1. **Detour on common knowledge.** Common knowledge might look like an innocuous assumption, but may have strong consequences in some situations. Consider the following story. Once upon a time, there was a village with 100 married couples. The women had to pass a logic exam before being allowed to marry. The high priestess was not required to take that exam, but it was common knowledge that she was truthful. The village was small, so everyone would be able to hear any shot fired in the village. The women would gossip about adulterous relationships and each knows which of the other husbands are unfaithful. However, no one would ever inform a woman that her husband is cheating on her.

The high priestess knows that not all husbands are faithful and decides that such immorality should not be tolerated. This is a successful religion and all women agree with the views of the priestess.

The priestess convenes all the women at the temple and *publicly announces* that the well-being of the village has been compromised—there is at least one cheating husband. She

Date: February 8, 2010.

also points out that even though none of them knows whether her husband is faithful, each woman knows about the other unfaithful husbands. She orders that each woman shoot her husband on the midnight of the day she finds out. 39 silent nights went by and on the 40^{th} shots were heard. How many husbands were shot? Were all the unfaithful husbands caught? How did some wives learn of their husband's infidelity after 39 nights in which *nothing* happened?

Since the priestess was truthful, there must have been at least one unfaithful husband in the village. How would events have evolved if there was exactly one unfaithful husband? His wife, upon hearing the priestess' statement and realizing that she does not know of any unfaithful husband, would have concluded that her own marriage must be the only adulterous one and would have shot her husband on the midnight of the first day. Clearly, there must have been more than one unfaithful husband. If there had been exactly two unfaithful husbands, then every cheated wife would have initially known of exactly one unfaithful husband, and after the first silent night would infer that there were exactly two cheaters and her husband is one of them. (Recall that the wives are all perfect logicians.) The unfaithful husbands would thus both be shot on the second night. As no shots were heard on the first two nights, all women concluded that there were at least three cheating husbands... Since shootings were heard on the 40^{th} night, it must be that exactly 40 husbands were unfaithful and they were all exposed and killed simultaneously.

For any measurable space X we denote by $\Delta(X)$ the set of probability measures (or distributions) on X.[1] A **mixed strategy** for player i is an element σ_i of $\Delta(S_i)$. A **correlated strategy profile** σ is an element of $\Delta(S)$. A strategy profile σ is **independent** (or **mixed**) if $\sigma \in \Delta(S_1) \times \ldots \times \Delta(S_n)$, in which case we write $\sigma = (\sigma_1, \ldots, \sigma_n)$ where $\sigma_i \in \Delta(S_i)$ denotes the marginal of σ on S_i. A **correlated belief** for player i is an element σ_{-i} of $\Delta(S_{-i})$. The set of **independent beliefs** for i is $\prod_{j \neq i} \Delta(S_j)$. It is assumed that player i has von Neumann-Morgenstern preferences over $\Delta(S)$ and u_i extends to $\Delta(S)$ as follows

$$u_i(\sigma) = \sum_{s \in S} \sigma(s) u_i(s).$$

[1]In most of our applications X is either finite or a subset of an Euclidean space.

2. Dominated Strategies

Are there obvious predictions about how a game should be played?

Example 1 (Prisoners' Dilemma). *Two persons are arrested for a crime, but there is not enough evidence to convict either of them. Police would like the accused to testify against each other. The prisoners are put in different cells, with no communication possibility. Each suspect is told that if he testifies against the other ("Defect"), he is released and given a reward provided the other does not testify ("Cooperate"). If neither testifies, both are released (with no reward). If both testify, then both go to prison, but still collect rewards for testifying. Each*

	C	D
C	$1, 1$	$-1, 2$
D	$2, -1$	$0, 0^*$

prisoner is better off defecting regardless of what the other does. Cooperation is a strictly dominated action for each prisoner. The only feasible outcome is (D, D), *which is Pareto dominated by* (C, C).

Example 2. *Consider the game obtained from the prisoners' dilemma by changing player 1's payoff for* (C, D) *from* -1 *to 1. No matter what player 1 does, player 2 still prefers*

	C	D
C	$1, 1$	$1, 2^*$
D	$2, -1$	$0, 0$

D *to* C. *If player 1 knows that 2 never plays* C, *then he prefers* C *to* D. *Unlike in the prisoners' dilemma example, we use an additional assumption to reach our prediction in this case: player 1 needs to deduce that player 2 never plays a dominated strategy.*

Formally, a strategy $s_i \in S_i$ is **strictly dominated** by $\sigma_i \in \Delta(S_i)$ if

$$u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}), \forall s_{-i} \in S_{-i}.$$

We can iteratively eliminate dominated strategies, under the assumption that "I know that you know that I know... that I know the payoffs and that you would never use a dominated strategy."

Definition 1. *For all $i \in N$, set $S_i^0 = S_i$ and define S_i^k recursively by*

$$S_i^k = \{s_i \in S_i^{k-1} |\ \nexists \sigma_i \in \Delta(S_i^{k-1}), u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}), \forall s_{-i} \in S_{-i}^{k-1}\}.$$

The set of pure strategies of player i that survive ***iterated deletion of strictly dominated strategies*** *is $S_i^\infty = \cap_{k \geq 0} S_i^k$. The set of surviving mixed strategies is*

$$\{\sigma_i \in \Delta(S_i^\infty) |\ \nexists \sigma_i' \in \Delta(S_i^\infty), u_i(\sigma_i', s_{-i}) > u_i(\sigma_i, s_{-i}), \forall s_{-i} \in S_{-i}^\infty\}.$$

Note that in a finite game the elimination procedure ends in a finite number of steps, so S^∞ is simply the set of surviving strategies at the last stage.

The definition above assumes that at each iteration all dominated strategies of each player are deleted simultaneously. Clearly, there are many other iterative procedures that can be used to eliminate strictly dominated strategies. However, the limit set S^∞ does not depend on the particular way deletion proceeds.[2] The intuition is that a strategy which is dominated at some stage is dominated at any later stage. Furthermore, the outcome does not change if we eliminate strictly dominated mixed strategies at every step. The reason is that a strategy is dominated against all pure strategies of the opponents if and only if it is dominated against all their mixed strategies. Eliminating mixed strategies for player i at any stage does not affect the set of strictly dominated pure strategies for any player $j \neq i$ at the next stage.

3. Rationalizability

Rationalizability is a solution concept introduced independently by Bernheim (1984) and Pearce (1984). Like iterated strict dominance, rationalizability derives restrictions on play from the assumptions that the payoffs and rationality of the players are common knowledge. Dominance: what actions should a player never use? Rationalizability: what strategies can a rational player choose? It is not rational for a player to choose a strategy that is not a best response to some beliefs about his opponents' strategies.

What is a "belief"? In Bernheim (1984) and Pearce (1984) each player i's beliefs σ_{-i} about the play of $j \neq i$ must be independent, i.e., $\sigma_{-i} \in \prod_{j \neq i} \Delta(S_j)$. Independent beliefs are consistent with the definition of mixed strategies, but in the context of an iterative procedure entail common knowledge of the fact that each player holds such beliefs. Alternatively, we

[2] This property does not hold for *weakly* dominated strategies.

may allow player i to believe that the actions of his opponents are correlated, i.e., any $\sigma_{-i} \in \Delta(S_{-i})$ is a possibility. The two definitions have different implications for $n \geq 3$. We focus on the case with correlated beliefs.

We can again iteratively develop restrictions imposed by common knowledge of the payoffs and rationality to obtain the definition of rationalizability.

Definition 2. *Set $S^0 = S$ and let S^k be given recursively by*

$$S_i^k = \{s_i \in S_i^{k-1} | \exists \sigma_{-i} \in \Delta(S_{-i}^{k-1}), u_i(s_i, \sigma_{-i}) \geq u_i(s_i', \sigma_{-i}), \forall s_i' \in S_i^{k-1}\}.$$

The set of ***correlated rationalizable strategies*** *for player i is $S_i^\infty = \cap_{k \geq 0} S_i^k$. A mixed strategy $\sigma_i \in \Delta(S_i)$ is rationalizable if there is a belief $\sigma_{-i} \in \Delta(S_{-i}^\infty)$ s.t. $u_i(\sigma_i, \sigma_{-i}) \geq u_i(s_i, \sigma_{-i})$ for all $s_i \in S_i^\infty$.*

The definition of **independent rationalizability** replaces $\Delta(S_{-i}^{k-1})$ and $\Delta(S_{-i}^\infty)$ above with $\prod_{j \neq i} \Delta(S_j^{k-1})$ and $\prod_{j \neq i} \Delta(S_j^\infty)$, respectively.

Definition 3. *A strategy $s_i^* \in S_i$ is a* ***best response*** *to a belief $\sigma_{-i} \in \Delta(S_{-i})$ if*

$$u_i(s_i^*, \sigma_{-i}) \geq u_i(s_i, \sigma_{-i}), \forall s_i \in S_i.$$

We say that a strategy s_i is **never a best response** for player i if it is not a best response to any $\sigma_{-i} \in \Delta(S_{-i})$. Recall that a strategy s_i of player i is **strictly dominated** if there exists $\sigma_i \in \Delta(S_i)$ s.t. $u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}), \forall s_{-i} \in S_{-i}$.

Theorem 1. *In a finite game, a strategy is never a best response if and only if it is strictly dominated.*

Proof. Clearly, a strategy s_i strictly dominated for player i by some σ_i cannot be a best response for any belief $\sigma_{-i} \in \Delta(S_{-i})$ as σ_i yields a strictly higher payoff than s_i against any such σ_{-i}.

We are left to show that a strategy which is never a best response must be strictly dominated. We prove that any strategy s_i of player i which is not strictly dominated must be a best response for some beliefs. Define the set of "dominated payoffs" for i by

$$D = \{x \in \mathbb{R}^{|S_{-i}|} | \exists \sigma_i \in \Delta(S_i), x \leq u_i(\sigma_i, \cdot)\}.$$

Clearly D is non-empty, closed and convex. Also, $u_i(s_i, \cdot)$ does not belong to the interior of D because it is not strictly dominated by any $\sigma_i \in \Delta(S_i)$. By the supporting hyperplane theorem, there exists $\alpha \in \mathbb{R}^{|S_{-i}|}$ different from the zero vector s.t. $\alpha \cdot u_i(s_i, \cdot) \geq \alpha \cdot x, \forall x \in D$. In particular, $\alpha \cdot u_i(s_i, \cdot) \geq \alpha \cdot u_i(\sigma_i, \cdot), \forall \sigma_i \in \Delta(S_i)$. Since D is not bounded from below, each component of α needs to be non-negative. We can normalize α so that its components sum to 1, in which case it can be interpreted as a belief in $\Delta(S_{-i})$ with the property that $u_i(s_i, \alpha) \geq u_i(\sigma_i, \alpha), \forall \sigma_i \in \Delta(S_i)$. Thus s_i is a best response to α. □

Corollary 1. *Correlated rationalizability and iterated strict dominance coincide.*

Theorem 2. *For every $k \geq 0$, each $s_i \in S_i^k$ is a best response (within S_i) to a belief in $\Delta(S_{-i}^k)$.*

Proof. Fix $s_i \in S_i^k$. We know that s_i is a best response within S_i^{k-1} to some $\sigma_{-i} \in \Delta(S_{-i}^{k-1})$. If s_i was not a best response within S_i to σ_{-i}, let s_i' be such a best response. Since s_i is a best response within S_i^{k-1} to σ_{-i}, and s_i' is a strictly better response than s_i to σ_{-i}, we need $s_i' \notin S_i^{k-1}$. This contradicts the fact that s_i' is a best response against σ_{-i}, which belongs to $\Delta(S_{-i}^{k-1})$. □

Corollary 2. *Each $s_i \in S_i^\infty$ is a best response (within S_i) to a belief in $\Delta(S_{-i}^\infty)$.*

Theorem 3. *S^∞ is the largest set $Z_1 \times \ldots \times Z_n$ with $Z_i \subset S_i, \forall i \in N$ s.t. each element in Z_i is a best response to a belief in $\Delta(Z_{-i})$ for all i.*

Proof. Clearly S^∞ has the stated property by Theorem 2. Suppose that there exists $Z_1 \times \ldots \times Z_n \not\subset S^\infty$ satisfying the property. Consider the smallest k for which there is an i such that $Z_i \not\subset S_i^k$. It must be that $k \geq 1$ and $Z_{-i} \subset S_{-i}^{k-1}$. By assumption, every element in Z_i is a best response to an element of $\Delta(Z_{-i}) \subset \Delta(S_{-i}^{k-1})$, contradicting $Z_i \not\subset S_i^k$. □

Example 3 (Rationalizability in Cournot duopoly)**.** *Two firms compete on the market for a divisible homogeneous good. Each firm $i = 1, 2$ has zero marginal cost and simultaneously decides to produce an amount of output $q_i \geq 0$. The resulting price is $p = \max(0, 1 - q_1 - q_2)$. Hence, if $q_1 + q_2 \leq 1$, the profit of firm i is given by $q_i(1 - q_1 - q_2)$. The best response correspondence of firm i is $B_i(q_j) = (1 - q_j)/2$ ($j = 3 - i$). If i knows that $q_j \lesseqqgtr q$ then $B_i(q_j) \gtreqqless (1 - q)/2$.*

We know that $q_i \geq q^0 = 0$ *for* $i = 1, 2$. *Hence* $q_i \leq q^1 = B_i(q^0) = (1 - q^0)/2$ *for all* i. *But then* $q_i \geq q^2 = B_i(q^1) = (1 - q^1)/2$ *for all* $i \ldots$ *We obtain*

$$\forall i,\ q^0 \leq q^2 \leq \ldots \leq q^{2k} \leq \ldots \leq q_i \leq \ldots \leq q^{2k+1} \leq \ldots \leq q^1,$$

where $q^{2k} = \sum_{l=1}^{k} 1/4^l = (1 - 1/4^k)/3$ *and* $q^{2k+1} = (1 - q^{2k})/2$ *for all* $k \geq 0$. *Clearly,* $\lim_{k \to \infty} q^k = 1/3$, *hence the only rationalizable strategy for firm* i *is* $q_i = 1/3$. *This is also the unique Nash equilibrium, which we define next.*

4. Nash Equilibrium

Many games are not solvable by iterated strict dominance or rationalizability. The concept of Nash (1950) equilibrium has more bite in some situations. The idea of Nash equilibrium was implicit in the particular examples of Cournot (1838) and Bertrand (1883) at an informal level.

Definition 4. *A mixed-strategy profile* σ^* *is a Nash equilibrium if for each* $i \in N$

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*), \forall s_i \in S_i.$$

Note that if a player uses a nondegenerate mixed strategy in a Nash equilibrium (one that places positive probability weight on more than one pure strategy) then he must be indifferent between all pure strategies in the support. Of course, the fact that there is no profitable deviation in pure strategies implies that there is no profitable deviation in mixed strategies either.

Example 4 (Matching Pennies)**.** *Pure strategy equilibria do not always exist.*

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

Nash equilibria are "consistent" predictions of how the game will be played—if all players predict that a specific Nash equilibrium will arise then no player has incentives to play differently. Each player must have correct "conjectures" about the strategies of their opponents and play a best response to his conjecture. Formally, Aumann and Brandenburger

(1995) provide a framework that can be used to examine the epistemic foundations of Nash equilibrium. The primitive of their model is an **interactive belief system** in which each player has a possible set of types, which correspond to beliefs about the types of the other players, a payoff for each action, and an action selection. Aumann and Brandenburger show that in a 2-player game, if the game being played (i.e., both payoff functions), the rationality of the players, and their conjectures are all *mutually known*, then the conjectures constitute a (mixed strategy) Nash equilibrium. Thus *common knowledge* plays no role in the 2-player case. However, for games with more than 2 players, we need to assume additionally that players have a common prior and and that conjectures are *commonly known*.

So far, we have motivated our solution concepts by presuming that players make predictions about their opponents' play by introspection and deduction, using knowledge of their opponents payoffs, knowledge that the opponents are rational, knowledge about this knowledge... Alternatively, we may assume that players extrapolate from past observations of play in "similar" games, with either current opponents or "similar" ones. They form expectations about future play based on past observations and adjust their actions to maximize their current payoffs with respect to these expectations. The idea of using learning-type adjustment processes originates with Cournot (1838). In that setting (Example 3), players take turns setting their outputs, each player choosing a best response to the opponent's last period action. Alternatively, we can assume simultaneous belief updating, best responding to sample average play, populations of players being anonymously matched, etc. If the process *converges* to a particular steady state, then the steady state is a Nash equilibrium. While convergence always occurs in Example 3, this is not always the case. How sensitive is the convergence to the initial state? If convergence obtains for all initial strategy profiles sufficiently close to the steady state, we say that the steady state is asymptotically stable. See figures 1.13-1.15 (pp. 24-26) in FT. The Shapley (1964) cycling example is interesting. Also, adjustment processes are myopic and do not represent a compelling description of behavior. Definitely such processes do not provide good predictions for behavior in the repeated game.

5. Existence and Continuity of Nash equilibria

Follow Muhamet's slides. We need the following result for future reference.

Theorem 4. *Suppose that each* S_i *is a convex and compact subset of an Euclidean space and that each* u_i *is continuous in* s *and quasi-concave in* s_i*. Then there exists a* ***pure strategy*** *Nash equilibrium.*