Game-Theoretic Artificial Intelligence[1] **Spring 2007**
Professor Amy Greenwald **Topic #2**

Matrix Games and Nash Equilibrium

This lecture is concerned with the Nobel Prize winning work of John Nash. In particular, we define the notion of mixed strategies in matrix games, and we present Nash's argument on the existence of mixed strategy (Nash) equilibrium.

1 Examples of Games

The most well-known game-theoretic scenario is the paradoxical situation known as the *Prisoners' Dilemma*, which was popularized by Axelrod [1] in his popular science book. The following is one (uncommon) variant of the story.[2]

A crime has been committed for which two prisoners are held incommunicado. The district attorney is assigned to question the prisoners. He designs the following incentive structure to induce the prisoners to talk. If neither prisoner talks, both prisoners automatically receive mild sentences (payoff 4). But if exactly one prisoner squeals on the other, the squealer is let off scot free (payoff 5), while the "squealee" is subject to a severe sentence (payoff 0). Finally, if both prisoners squeal, they share the severe punishment (payoff 1).

The Prisoners' Dilemma is a two player a matrix—or strategic, or normal form—game. Such games are easily described by payoff matrices, where the strategies of player 1 and player 2 serve as row and column labels, respectively, and the corresponding payoffs are listed as pairs in matrix cells such that the first (second) number is the payoff to player 1 (2). The payoff matrix which describes the Prisoners' Dilemma is depicted in Figure 1, with C denoting "cooperate" or "confess", and D denoting "defect" or "don't cooperate."

This game is known as the Prisoners' Dilemma because the only rational outcome is (D, D), which yields suboptimal payoffs of $(1, 1)$. The reasoning is as follows. If player 1 plays C, then player 2 is better off playing D, since D yields a payoff of 5, whereas C yields only 4; but if player 1 plays D, then player 2 is again better off playing D, since D yields a payoff of 1, whereas C yields only 0. Hence, regardless of the strategy of player 1, a rational player 2 plays D. By a symmetric argument, a rational player 1 also plays D. Thus, the unique outcome of the game, assuming the players are rational, is (D, D).

[2]The original anecdote due to A.W. Tucker appears in Rapoport [5]; the latter author is the two-time winner of the Prisoners' Dilemma computer tournament organized by Axelrod.

1 \ 2	*C*	*D*
C	4,4	0,5
D	5,0	1,1

Figure 1: The Prisoners' Dilemma

Another popular two-player game is called the *Battle of the Sexes*. A man and a woman would like to spend an evening out together; however, the man prefers to go to a football game (strategy F), while the woman prefers to go to the ballet (strategy B). Both the man and the woman prefer to be together, even at the event that is not to their liking, rather than go out alone. The payoffs of this coordination game are shown in Figure 2; the woman is player 1 and the man is player 2. In this game, there are two coordination equilibria, one which is preferred by the woman, and another which is preferred by the man.

1 \ 2	*B*	*F*
B	2,1	0,0
F	0,0	1,2

Figure 2: Battle of the Sexes

The stag hunt game (see Figure 3) is a prototypical social contract. Rousseau tells an early version of the story in *A Discourse on Inequality.*:

> If it was a matter of hunting a deer, everyone realized that he must remain well faithful to his post; but if a hare happened to pass within reach of one of them, we cannot doubt that he would have gone off in pursuit of it without scruple...

This game has two equilibria, one of which Pareto dominates the other: *i.e.*, all players are simultaneously better off at one equilibrium than the other. But action H *risk-dominates* action D, since action H is safer for each player, given his/her uncertainty about the other player's action.

1 \ 2	*D*	*H*
D	2,2	0,1
H	1,0	1+ε,1+ε

Figure 3: Stag Hunt

Our final example is an ecological game which was studied by Maynard Smith [6] in his analysis of the theory of evolution in terms of games (see Figure 4). The game is played between animals of the same size who live in the wilderness and encounter one another in their search for prey. During an encounter between two animals, each animal has a choice between behaving as a hawk: *i.e.*, fighting for the prey; or as a dove: *i.e.*, sharing the prey peacefully. If both animals decide to play like hawks, then each animal has an equal chance of winning the value v of the prey or of losing the fight at cost c, where $0 < v < c$; thus, the expected payoff to both players is $(v - c)/2$. Alternatively, if both animals act as doves, then the prey is shared with equal payoffs $v/2$. Finally, if one animal behaves like a hawk and the other behaves like a dove, then the hawk gets a payoff worth the full value of the prey and the other gets nothing. In this game, the animals prefer to choose opposing strategies: if one animal plays hawk, then it is in the best interest of the other to play dove; and inversely, if one animal plays dove, then it is in the best interest of the other to play hawk.

1 \ 2	*H*	*D*
H	(v-c)/2, (v-c)/2	v,0
D	0,v	v/2,v/2

Figure 4: Hawks and Doves

2 Nash Equilibrium

A *Nash equilibrium* is a strategy profile from which none of the players has any incentive to deviate. In particular, no player can achieve strictly greater payoffs by choosing any strategy other than the one prescribed by the profile, given that all other players choose their prescribed strategies. In this sense, a Nash equilibrium specifies optimal strategic choices for all players.

Let us examine the Nash equilibria in the aforementioned examples:

- In the Prisoners' Dilemma, (D, D) is a Nash equilibrium: given that player 1 plays D, the best response of player 2 is to play D; given that player 2 plays D, the best response of player 1 is to play D.
- The Battle of the Sexes has two pure strategy Nash equilibria, namely (B, B), and (F, F): If the woman plays B, then the best response of the man is B; if the man plays B, then the best response of the woman is B. Analogously, if the woman plays F, the best response of the man is F; if the man plays F, the best response of the woman is F.
- The game of Hawks and Doves also has two pure strategy Nash equilibria, only this time, the players seek to miscoordinate their actions, rather than coordinate. If one player plays the hawk, then the other player prefers to play the dove, achieving payoffs of 0, rather than $(v - c)/2 < 0$. On the other hand, if one player plays the dove, then the other player prefers to play the hawk, achieving payoffs of v, rather than $v/2$.
- Finally, in the Stag Hunt, both (D, D) and (H, H) are Nash equilibria: hunting for deer is Pareto-optimal, but hunting for hare is risk-dominant.

Matching Pennies is another well-known example of a two player, zero-sum game. In this game, each of the players, the *matcher* and the *mismatcher*,[3] flips a coin, and the payoffs are determined as follows. If the coins come up matching (*i.e.*, both heads or both tails), then the matcher wins, so the mismatcher pays the matcher \$1. If the coins do not match (*i.e.*, one head and one tail), then the mismatcher wins, so the matcher pays the mismatcher \$1. In Figure 5, player 1 is the mismatcher and player 2 is the matcher. This game is called zero-sum because the payoffs in each cell of the matrix sum to zero.

In the game of Matching Pennies, there is no pure strategy Nash equilibrium. If player 1 plays H, then the best response of player 2 is T; but if player 2 plays T, the best response of player 1 is not H, but T. Moreover, if player 1 plays T, then the best response of player 2 is H; but if player 2 plays H, then the best response of player 1 is not T, but H. This game, however, does have a mixed strategy Nash equilibrium. A mixed strategy is a randomization over a set of

[3]The mismatcher is often affectionately referred to as Miss Matcher.

1 \ 2	H	T
H	1,-1	-1,1
T	-1,1	1,-1

Figure 5: Hawks and Doves

pure strategies. In particular, the probabilistic strategy profile in which both players choose H with probability $\frac{1}{2}$ and T with probability $\frac{1}{2}$ is the unique (mixed strategy) Nash equilibrium in the game of Matching Pennies.

A *matrix game* is a 3-tuple $\Gamma = (N, (A_i, R_i)_{1 \leq i \leq n})$, where

- N is a set of n players
- A_i is a finite strategy (or action) set ($a_i \in A_i$)
- $R_i : A \to \mathbb{R}$ is a payoff function, where $A = A_1 \times \ldots \times A_n$

Matrix games are also sometimes called games in *strategic*, or *normal*, form.

In this formalism, the Prisoners' Dilemma consists of a set of players $N = \{1, 2\}$, with strategy (action) sets $A_1 = A_2 = \{C, D\}$, and payoffs as follows:

$$\begin{array}{ll} R_1(C, C) = R_2(C, C) = 4 & R_1(C, D) = R_2(D, C) = 0 \\ R_1(D, D) = R_2(D, D) = 1 & R_1(D, C) = R_2(C, D) = 5 \end{array}$$

A mixed strategy set for player i is the set of probability distributions over the action set A_i, which can be described by the simplex operator Δ:

$$\Delta(A_i) = \left\{ q_i : A_i \to [0, 1] \;\; | \sum_{a_i \in A_i} q_i(a_i) = 1 \right\}$$

For convenience, let $Q_i \equiv \Delta(A_i)$. The usual notational conventions extend to mixed strategies: *e.g.*, $Q = \prod_i Q_i$ and $q = (q_i, q_{-i}) \in Q$. In the context of mixed strategies, the expected payoffs to player i from strategy profile q are:

$$\mathbb{E}_{a \sim q}[R_i(a)] = \sum_{a \in A} q(a) R_i(a)$$

where

$$q(a) = \prod_{j=1}^{N} q_j(a_j)$$

As usual, the payoffs to player i depend on the mixed strategies of all players.

An implication of the assumption of rationality is that a rational player always plays an optimizing strategy, or a *best response* to the strategies of the other players. A strategy $q_i^* \in Q_i$ is a *best response* for player i to opposing strategy $q_{-i} \in Q_{-i}$ iff $\forall q_i \in Q_i$, $\mathbb{E}[R_i(q_i^*, q_{-i})] \geq \mathbb{E}[R_i(q_i, q_{-i})]$. The best response set for player i to strategy profile q_{-i} is:

$$\text{BR}_i(q_{-i}) = \{q_i^* \in Q_i \mid \forall q_i \in Q_i,\ \mathbb{E}[R_i(q_i^*, q_{-i})] \geq \mathbb{E}[R_i(q_i, q_{-i})]\}$$

The set $\text{BR}_i(q_{-i})$ is often abbreviated $\text{BR}_i(q)$. Let $\text{BR}(q) = \prod_i \text{BR}_i(q)$.

A Nash equilibrium is a strategy profile in which all players choose strategies that are best responses to the strategic choices of the other players. Characterized in terms of best response sets, a Nash equilibrium is a strategy profile q^* *s.t.* $q^* \in \text{BR}(q^*)$. It is apparent from this definition that a Nash equilibrium is a fixed point of the best response relation. In 1951, Nash proved that *every finite strategic form game has (at least) one mixed strategy Nash equilibrium* [4]. The existence proof relies on a fundamental result in topology, namely Kakutani's fixed point theorem, a generalization of Brouwer's fixed point theorem [2].

2.1 Digression: Fixed Point Theorems

A *closed* set X contains all its limit points: *i.e.*, for all sequences $\{x^n\}$ in X, if $x^n \to x$, then $x \in X$. A set is *open* if its complement is closed. For example, the interval $[0,1]$ is a closed subset of $\mathbb{R}$ and the interval $(0,1)$ is an open subset of $\mathbb{R}$; but the intervals $(0,1]$ and $[0,1)$ are neither open nor closed subsets of $\mathbb{R}$.

A *subsequence* of the sequence $\{x^n\}$ is a restriction of the mapping from $\mathbb{N}$ to X to a mapping from $N' \subseteq \mathbb{N}$ to X. A subsequence is a sequence in its own right. A set X is called *compact* iff every sequence in X has a convergent subsequence that converges to a point in X. Note that every compact set is closed.

Heine-Borel Theorem A set $X \subseteq \mathbb{R}^n$ is compact iff it is closed and bounded.

Weierstrass' Theorem A real-valued continuous function on a compact set attains a maximum and a minimum.

A *correspondence* $f : X \Rightarrow Y$ is a mapping from X to subsets of Y: *i.e.*, $f(x) \subseteq Y$. A function is a special case of a correspondence in which $f(x)$ is a singleton for all $x \in X$. Just as functions can be associated with their graphs, correspondences can be associated with their graphs. The *graph* of the correspondence f is the set of pairs $\{(x, y) \mid x \in X, y \in f(x)\}$. A set X is called *convex* iff $\forall \lambda \in [0,1]$, $\forall x, y \in X$, $\lambda x + (1-\lambda)y \in X$. The correspondence f is *nonempty* and *convex-valued* iff $f(x)$ is nonempty and convex for all $x \in X$.

Kakutani's Fixed Point Theorem. Let $X \subset \mathbb{R}^n$ be nonempty, compact, and convex. If $f : X \Rightarrow X$ is a nonempty, convex-valued correspondence with a closed graph, then f has a fixed point: *i.e.*, there exists $x^* \in X$ *s.t.* $x^* \in f(x^*)$.

Brouwer's Fixed Point Theorem. Let $X \subset \mathbb{R}^n$ be nonempty, compact, and convex. If $f : X \to X$ is a continuous function, then f has a fixed point: *i.e.*, there exists $x^* \in X$ *s.t.* $x^* = f(x^*)$.

2.2 Proof of Existence

The proof of existence of Nash equilibrium is a direct application of Kakutani's fixed point theorem. It suffices to show that the set of mixed strategies Q is nonempty, compact, and convex, and that the best-response correspondence (*i.e.*, $\text{BR} : Q \rightrightarrows Q$) is nonempty and convex-valued, with a closed graph.

Lemma The set of mixed strategies Q is nonempty, compact, and convex.

Proof Recall that $Q = \prod_i Q_i$, where Q_i is the set of probability over distributions player i's the action set A_i. The set Q is nonempty (assuming A_i is nonempty, for all players i).

Given a sequence $\{(q_1^m, \ldots, q_n^m)\}$ of mixed strategies. that converges to $(q_1^*, \ldots, q_n^*)$. This limit point is indeed a mixed strategy: *i.e.*, $q_i^* \geq 0$ and $\sum_i q_i^* = 1$. The former claim follows from the fact that the limit of a sequence of non-negative points is itself non-negative. The latter claim follows from the fact that the sum of the limits equals the limit of the sum. Thus, Q is closed. Moreover, Q is bounded in each component by 0 and 1. Therefore, Q is compact.

The set of mixed strategies Q_i for each player i is convex: *i.e.*, for all $q_i, p_i \in Q_i$, for all $\lambda \in [0,1]$, the convex combination $\lambda q_i + (1-\lambda)p_i \in Q_i$. thus, given two elements $(q_1, \ldots, q_n), (p_1, \ldots, p_n) \in Q$, the convex combination $\lambda(q_1, \ldots, q_n) + (1-\lambda)(p_1, \ldots, p_n) = (\lambda q_1 + (1-\lambda)p_1, \ldots, \lambda q_n + (1-\lambda)p_n) \in Q$, for all $\lambda \in [0,1]$.

Thus, the set of mixed strategies Q is nonempty, compact, and convex. □

Lemma The best-response correspondence is nonempty.

Proof By Weierstrass' theorem, any real-valued continuous function on a compact set attains a maximum. Recall that the set Q_i is compact. Since R_i is a linear function of Q_i, R_i is continuous. Thus, $\text{BR}_i : Q \to Q_i$ is nonempty, for all players i, from which it follows that BR is nonempty. □

Lemma The best-response correspondence is convex-valued.

Proof If $q_i^*, p_i^* \in \text{BR}_i(q_{-i})$ are best replies of player i to q_{-i}, then $R_i(q_i^*, q_{-i}) = R_i(p_i^*, q_{-i}) = \lambda R_i(q_i^*, q_{-i}^*) + (1-\lambda)R_i(p_i^*, q_{-i}^*)$. Now, by the linearity of R_i, $\lambda R_i(q_i^*, q_{-i}^*) + (1-\lambda)R_i(p_i^*, q_{-i}^*) = R_i(\lambda q_i^* + (1-\lambda)p_i^*, q_{-i}^*)$. Thus, the convex combination $\lambda q_i^* + (1-\lambda)p_i^* \in \text{BR}_i(q_{-i})$. Since q_{-i} was arbitrary, BR_i is convex-valued. Since i was arbitrary, BR is convex-valued. □

Lemma The graph of the best-response correspondence is closed.

Proof Must show $p \in \textsc{br}(q)$, given the sequences $q^m, p^m \in Q$ *s.t.* $q^m \to q$ and $p^m \to p$, with $p^m \in \textsc{br}(q^m)$ for all m. Suppose not: *i.e.*, suppose there exists player i *s.t.* $p_i \notin \textsc{br}_i(q_{-i})$. It follows that there exists $q_i \in Q_i$ *s.t.* $R_i(q_i, q_{-i}) > R_i(p_i, q_{-i})$. Now let $\delta \equiv R_i(q_i, q_{-i}) - R_i(p_i, q_{-i}) > 0$. Since R_i is linear, and therefore continuous, for all $\epsilon > 0$, there exists $M_\epsilon \in \mathbb{N}$ *s.t.* for all $m \geq M_\epsilon$, $|R_i(p_i^m, q_{-i}^m) - R_i(p_i, q_{-i})| < \epsilon$ and $|R_i(q_i, q_{-i}^m) - R_i(q_i, q_{-i})| < \epsilon$. Now

$$\begin{aligned} R_i(q_i, q_{-i}^m) &> R_i(q_i, q_{-i}) - \epsilon \\ &= R_i(p_i, q_{-i}) + R_i(q_i, q_{-i}) - R_i(p_i, q_{-i}) - \epsilon \\ &= R_i(p_i, q_{-i}) + \delta - \epsilon \\ &> R_i(p_i^m, q_{-i}^m) + \delta - 2\epsilon \end{aligned}$$

If $\epsilon = \delta/2$, then $R_i(q_i, q_{-i}^m) > R_i(p_i^m, q_{-i}^m)$, for all $m \geq M_{\delta/2}$. But then $p_i^m \notin \textsc{br}_i(q_{-i}^m)$ for all m. Contradiction. Therefore, the graph of $\textsc{br}$ is closed. □

Exercise Compute the best-response correspondences for the game depicted in Figure 6, a version of Hawks and Doves. Plot these correspondences, and compute all Nash equilibria.

1 \ 2	*H*	*D*
H	−1,−1	2,0
D	0,2	1,1

Figure 6: Hawks and Doves

3 Summary

In this lecture, we defined Nash equilibrium and reproved Nash's theorem guaranteeing its existence in all (finite) matrix games.

Although Nash equilibrium is the generally accepted solution concept in the deductive analysis of matrix games, the Nash equilibria in our examples are somewhat peculiar. In the Prisoners' Dilemma, the Nash equilibrium payoffs are sub-optimal. In the game of Matching Pennies, there is no pure strategy Nash equilibrium; the unique Nash equilibrium is probabilistic. Finally, in the coordination and miscoordination games, the Nash equilibrium is not unique.

In future lectures, alternative notions of equilibria are discussed.

References

[1] R. Axelrod. *The Evolution of Cooperation.* Basic Books, New York, 1984.

[2] L. Brouwer. Uber abbildungen von mannigfaltigkeiten. *Mathematische Annalen*, 38(71):97–115, 1912.

[3] S. Kakutani. A generalization of brouwer's fixed point theorem. *Duke Mathematical Journal*, 8:457–458, 1941.

[4] J. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.

[5] A. Rapoport. *Two-Person Game Theory.* University of Michigan Press, Ann Arbor, 1966.

[6] J.M. Smith. *Evolution and the Theory of Games.* Cambridge University Press, New York, 1982.