MIT 14.11

Pure Strategy Matrix Form Games and Nash Equilibria

Zoë Hitzig, Moshe Hoffman, and Erez Yoeli

September 18, 2013

1   Introduction

Game theory models situations where multiple "players" (firms, states, animals, people), play strategies (e.g. sacrifice to help another, grab for a contested object, mate with), and then receive "payoffs" (dollars, offspring). The payoffs to each depend on their strategies as well as the strategys of the other players.

In this note we will introduce the most basic type of "game"–**pure-strategy matrix-form games**. Such games have only 2 players, each player only has finitely many possible strategies, and each strategy is determined in advance of play.

Pure-strategy matrix-form games can be represented concisely in a table as in the figure below, known as a "payoff matrix." These matrices can be read as follows: Each row is labeled with a possible strategy for player 1 (A and B). Each column is labeled with a possible strategy for player 2 (C and D). For any given row and column, the first number within the corresponding cell refers to the payoffs player 1 gets when player 1 plays that row and player 2 plays that column (e.g. row A column D gives player 1 c). The second number within the corresponding cell refers to the payoffs player 2 gets when player 1 plays that row and player 2 plays that column (e.g. row A column D gives player 2 d).

	C	D
A	a, b	c, d
B	e, f	g, h

Pure-strategy matrix-form games include the familiar, prisoners' dilemma, coordination game, and hawk-dove game, which we will use to analyze social behavior in upcoming lectures.

The prisoners' dilema models cooperation, where each player simultaneously play cooperate or defect, where both players do better if they both cooperate than if they both defect, but each does better by defecting, holding constant the other player's strategy, as in:

	Cooperate	Defect
Cooperate	r, r	s,t
Defect	t, s	$\mathbf{p,p}$

$$t > r > p > s$$

The coordination game models, umm, coordination, where each player simultaneously plays a strategy (say A or B), and they get higher payoffs if they play the same strategy as the other, as in:

	A	B
A	$\mathbf{a,a}$	b,c
B	c,b	$\mathbf{d,d}$

$$a > c, d > b$$

Hawk-Dove models conflict over a contested item, with value v, where each player precommits to behave either aggressively ("hawk") or not (dove). If only one player plays aggressively, he gets the contested item. If neither behave aggressively, they are equally likely to obtain the item. If they both behave aggressively, they pay the cost c of conflict, and again are equally as likely to obtain the item. It is presumed that the cost of conflict is sufficiently high that it isn't worth playing aggressively if the other is going to play aggressively.

	Hawk	Dove
Hawk	$\frac{v}{2} - c, \frac{v}{2} - c$	$\mathbf{v,0}$
Dove	$\mathbf{0,v}$	$\frac{v}{2}, \frac{v}{2}$

$$c > v/2 > 0$$

We will later use these simple games to gain insight into prosocial behavior, and less obviously, rights, the ommission commission distinction, norms against chemical weapons, and other puzzling social behaviors.

We will also introduce the most basic tool game theorists use to "solve" such games–**pure Nash equilibrium**. Intuitively, a pure Nash equilibrium is a specification of a strategy for each player such that no player would benefit by changing his strategy, provided the other players don't change their strategies.

This concept, as simple as it sounds, often leads to counterintuitive "solutions" (bolded in above figures). In the prisoners' dilemma, we will show, mutual defection is the only Nash equilibria. This is counterintuitive because each would do better if both cooperated than if both defect. This begs the question of why do players ever behave prosocially, which we will address in later lectures. In the coordination game, there are two equilibria, one where both play A and one where both play B. This is also counterintuitive because, both playing B gives lower payoffs than both playing A. Lastly, in the hawk-dove game, there are two equilibria, one where player 1 plays hawk, and player 2 plays dove, and one where player 2 plays hawk and player 1 plays dove. This is again counterintiuitive, because in the first equilibrium player 1 does much better than player 2, even though the two players have exactly identical payoffs and strategys available (and conversely in the second equilibrium).

Moreover, we will later see that evolutionary and learning processes prevent populations from "stabilizing" at strtegies that are not pure Nash equilibrium, making this "solution" a very good starting point for gaining insights into the above social puzzles. For instance, evolution and learning will lead to mutual defection in the prisoner's dilema, may lead to inefficient coordination in the coordination game, and will lead to "arbitrary" assignment of "property rights" in the hawk-dove game. Each of these predictions will help us explain some social puzzles and lead to some novel predictions, or powerful prescriptions.

Now for some formalism. (See math handout titled "Math Preliminaries" for basic math background which will be assumed throughout. Also note that this handout is adapted from Martin J. Osborne's *Introduction to Game Theory* Chapter 2.)

2 Pure Strategy Games

<u>Definition 1.1</u>: A **pure-strategy matrix-form game** is a two player game with finite strategies, specified by four pieces of information: the set of strategies for each player, and a function, again for each player, that defines preferences over each pair of strategies. Formally, we can define any

game as an ordered pair $< \{S_i\}_{i=1,2}, \{U_i\}_{i=1,2} >$ where:

- $\{S_i\}_{i=1,2}$ is the finite set of strategies available to each player.

  E.g.: $i = 1, 2$, $S_i =$ {Cooperate, Defect}.

    - $s_i \in S_i$ is a generic strategy available to player i, i.e., cooperate.
    - $S = \prod_{i=1,2} S_i$ is the set of strategy "profiles," i.e. the set of possible pairs of strategies, one for each player.

      E.g.: S = {(Cooperate, Cooperate), (Cooperate, Defect), (Defect, Cooperate), (Defect, Defect)}.

    - $s = (s_1, s_2) \in S$ is a generic element of S.

      E.g.: $s =$(Cooperate, Defect). It is common to refer to s as (s_i, s_{-i}), where s_{-i} represents the strategies of everyone other than i.

- U_i is a function from $S \to \mathbb{R}$ that defines i's preferences over all strategy profiles.

    - E.g.: U_1(Cooperate, Defect) = s while U_1(Defect, Defect) = p. Notice that the lowercase s to the left is a payoff pulled from the matrix representing the PD game, and not a strategy.
    - It is common to say "Player i prefers s over s' iff $U_i(s) \geq U_i(s')$." This is technically ambiguous since it doesn't specify others' actions. When this is said, others' actions are being held fixed, as in, "Player 1 prefers (Cooperate, Defect) to (Defect, Defect)."

Notice how precisely a payoff matrix can represent the above information. That's why, usually, payoffs are represented in the matrix instead of specifying $< \{S_i\}_{i=1,2}, \{U_i\}_{i=1,2} >$ explicitly.

Now that we've laid out this notation, we're ready to define a Nash equilibrium.

<u>Definition 1.2</u>: A strategy profile $s \in S$ is a **pure Nash equilibrium** if $\forall i$ and $\forall s_i \in S_i$, $U_i(s_i, s_{-i}) \geq U_i(s'_i, s_{-i})$.

In plain terms, a pure Nash equilibrium is a strategy profile in which no player would benefit by deviating, given that all other players don't deviate. Some games have multiple pure Nash equilibria and some games do not have any pure Nash equilibria.

Pure Nash equilibria are easy to find in matrix form games, using what we will call the "starring algorithm." As we will show in the following examples. The algorithm works as follows: For each

column, add an asterisk superscript to the highest number corresponding to player 1's payoff in that column. For each row, add an asterisk superscript to the highest number corresponding to player 2's payoff in that row. Any cell that has two stars is a pure Nash equilibrium. A formal solution, in terms of the utility function, is laid out in the Appendix.

	Cooperate	Defect
Cooperate	r, r	$s, t*$
Defect	$t*, s,$	$p*, p*$

$t > r > p > s$

	A	B
A	$a*, a*$	b, c
B	c, b	$d*, d*$

$a > c, d > b$

	Hawk	Dove
Hawk	$\frac{v}{2} - c, \frac{v}{2} - c$	$v*, 0*$
Dove	$0*, v*$	$\frac{v}{2}, \frac{v}{2}$

$c > v/2 > 0$

3   Extensions and Further Reading

Our formal definition of matrix-form games extends easily to more general definitions in a few important ways. If more than two players are allowed (i.e. $N > 2$), we can no longer represent the entire game in a matrix. With $N > 2$ we represent a game as a higher dimensional array. Our definition of matrix-form games also requires that the number of available strategies is finite. A normal-form game allows for any finite number of players N as well as an infinite number of strategies available to each player. The definition of a normal-form game is therefore a straightforward extension, as is the definition of pure Nash equilibrium in normal-form games, but the algorithm used for finding pure Nash equilibrium in normal form games is different from the starring algorithm. These extensions are typically covered in introductory game theory classes, but are not used in any of our social applications. For discussion of these extensions, see Osborne.

4 Appendix

Let's assign concrete values to the generic Prisoner's Dilemma outlined above, and compare each players' preferences over each strategy pairwise, finding the Nash equilibrium in the most thorough way possible.

$$
\begin{array}{ccc}
 & C & D \\
C & -1, -1 & -4, 0 \\
D & 0, -4 & -3, -3
\end{array}
$$

U_1(Cooperate, Cooperate) $= -1$ U_1(Cooperate, Defect)$= -4$

U_2(Cooperate, Cooperate)$= -1$ U_2(Cooperate, Defect)$= 0$

U_1(Defect, Cooperate)$= 0$ U_1(Defect, Defect)$= -3$

U_2(Defect, Cooperate)$= -4$ U_2(Defect, Defect)$= -3$

In order to find the pure Nash equilibria, we must test each strategy profile and compare all payoffs above pairwise.

- Is (Cooperate, Cooperate) a Nash equilibrium?

  U_1(Cooperate, Cooperate)$= -1 \not\geq 0 = U_1$(Defect, Cooperate)

  U_2(Cooperate, Cooperate)$= -1 \geq 0 = U_2$(Cooperate, Defect)

  No.

- Is (Cooperate, Defect) a Nash equilibrium?

  U_1(Cooperate, Defect)$= -4 \not\geq -3 = U_1$(Defect, Defect)

  U_2(Cooperate, Defect)$= 0 \geq -1 = U_2$(Cooperate, Cooperate)

  No.

- Is (Defect, Cooperate) a Nash equilibrium?

  U_1(Defect, Cooperate)$= 0 \geq -3 = U_1$(Defect, Defect)

  U_2(Defect, Cooperate)$= -4 \not\geq -3 = U_2$(Cooperate, Cooperate)

  No.

- Is (Defect, Defect) a Nash equilibrium?

  U_1(Defect, Defect)$= -3 \geq -4 = U_1$(Cooperate, Defect)

  U_2(Defect, Defect)$= -3 \geq -4 = U_2$(Defect, Cooperate)

Yes, because both of the above inequalities hold.

Of course it would be far less cumbersome to use the algorithm described at the end of the previous section to test each Nash equilibrium.

14.11 Insights from Game Theory into Social Behavior
Fall 2013