1 Stable equilibria

Today we will consider the stability of equilibria. Natural game play does not reach Nash equilibria. The artificial intelligence community has looked into good strategies for game play, and there has been interest in learning good strategies, which we will cover in the second half of this lecture.

The first question we will consider is this: if players are in an equilibrium, and one player changes his/her strategy, what happens to the equilibrium? Consider for example the *Battle of Sexes* presented in the first lecture:

	B	S
B	$2,1$	$0,0$
S	$0,0$	$1,2$

We noted that (B,B) and (S,S) were the pure Nash equilibria. But suppose that two players were in the (B,B) Nash equilibrium and one player decides to choose S. The other player must now choose S as its Nash strategy. The players are once again in a Nash equilibrium! However, for the *prisoner's dilemma*:

	C	D
C	$3,3$	$0,4$
D	$4,0$	$1,1$

The unique pure Nash equilibrium is (D,D). If one player changes his/her strategy to C, then the other player does not have any incentive to change their strategy.

We define an equilibrium to be strict if for all players their best response is unique. So the *Battle of Sexes* does not have a strict equilibrium, but the *prisoner's dilemma* does.

Now we consider modeling the evolution of populations in biology. There is a natural application of games to evolution. We can consider the payoff of a strategy as the reproductive success of a player.

For biological applications, it makes sense to consider symmetric games, e.g. games in which the payoff matrix is the same for all players. In 2 player games, these are represented by symmetric matrices as above in the *prisoner's dilemma*. The first game we will consider is *Dove & Hawk*. Here is our first version:

	D	H
D	$1/2, 1/2$	$0,1$
H	$1,0$	$1/4, 1/4$

The way it works is as follows. Suppose there is some food source, and player A and player B arrive at the food source. If two doves arrive, they peaceably share the food. If a dove and a hawk arrive, naturally the hawk gets the food. Finally, if two hawks arrive, they fight. Their reward will be less than the two doves. In this version, a pure Nash equilibrium is (H,H). In fact this is the strict equilibrium.

So the question concering biologists is: given a population consisting of p proportion of doves and 1-p hawks, is it better to play dove or hawk? Note that by changing the 1/4,1/4 in the payoff matrix we can change the social optimum or even what the Nash equilibria are. One might imagine in some payoff matrices, the best strategy depends on the proportion p.

We define a strategy to be an Evolutionary Stable Strategy (ESS) as: if everyone plays stable strategies and ε percentage plays some other strategy, then there is a small enough $\varepsilon > 0$ so that those playing the this alternate strategy get a smaller reward.

Consider the following variation of the *Dove & Hawk* game above:

	D	H
D	$1/2, 1/2$	$0, 1$
H	$1, 0$	$0, 0$

Now (H,H) is still a Nash equilibrium. But it is not a strict equilibrium, since there is no penalty for playing dove when the other player plays H. It is an ESS though: Suppose $\varepsilon > 0$ plays D

- for doves D
  - $1 - \varepsilon$ meet H, get a reward of 0
  - ε meet D, get a reward of 1/2
- for hawks H
  - $1 - \varepsilon$ meet H, get a reward of 0
  - ε meet D, get a reward of 1

Now consider the following *Dove & Hawk* game:

	D	H
D	$1/2, 1/2$	$0, 1$
H	$1, 0$	$-1/2, -1/2$

The pure Nash equilibria are (H,D) (D,H). Mixing probabilities (1/2, 1/2) is also an equilibrium. This mixed equilibrium is evolutionarily stable. Thus a population of 1/2 hawks and 1/2 doves is stable.

Now we turn our attention to learning. Consider a game based on *prisoner's dilemma*. Two players meet in a match, and each chooses to either cooperate or defect. After they have chosen, the choices for each player are revealed, and the payoff for each player is calculated and added to each player's score. The players retain knowledge of prior games against this opponent. The artificial intelligence community has played a version of this game where there are a set of n players, and each player plays each other player k times, for some k. The goal is to maximize the sum of the rewards from each match.

Consider the following strategies. Suppose we knew that the other player would always defect. Then optimally we would also always choose to defect. On the other hand, if the other player always cooperated, we would optimally always choose to defect. So an obvious strategy is to always defect. But of course, if everyone does this then we score fewer points than if everyone had chosen to cooperate.

Consider another strategy, *tit-for-tat*. In this strategy, for each time an opponent defects, we defect once. Thus if we are paired with a pure defector, we only come out 1 point the worse over k matches. If we are paired with a pure cooperator, we do as well as a pure cooperation strategy. If we are paired with another *tit-for-tat* player, we also do as well as a pure cooperation strategy. Note that in the last strategy, we do better than if we had defected always, and in the first strategy we do better than if we had cooperated always. So *tit-for-tat* would seem to be a good strategy.

Suppose we had a population consisting of C, D and *tit-for-tat*. In the evolution game a population of *tit-for-tat* results. We can also show that tit-for-tat is an ESS, while neither C nor D is an ESS.

But if there are other variants, for example *tit-for-2-tat*, then some other strategy may win. In the games that the AI people play, other variants have shown superior to *tit-for-tat*.

Next week we will look at market equilibrium.