CAP 5993/CAP 4993
Game Theory

Instructor: Sam Ganzfried

sganzfri@cis.fiu.edu

Schedule

- HW4 out 4/4 due 4/13 (HW2 and HW3 back by 4/4).
- Project presentations on 4/18 and 4/20.
- Project writeup due 4/20.
- Final exam on 4/25.

Projects

- Can work in groups 1-3
- Project can be theoretical, or applied
  - Could involve implementation, e.g., with Gambit
- Original summary project is ok if it is approved by me
- Can get full credit for all project types

Solution concepts

- …

Solution concepts

- Maxmin strategies
- Weak/strict domination
- Nash equilibrium
- Refinements of Nash equilibrium
  - Trembling hand perfect equilibrium
  - Subgame perfect equilibrium
  - Proper equilibrium
  - Evolutionarily stable strategies
- Quantal response equilibrium

Game representations

- …

Game representations

- Strategic form
- Extensive form
  - Perfect information
  - Perfect information (with chance events)
  - Imperfect information (with chance events)
- Repeated (finitely and infinitely)

Battle of the sexes

	F	**C**
F	2,1	0,0
C	0, 0	1,2

- 3 equilibria:
  - (F,F) (the payoff is (2,1))
  - (C,C) (payoff is (1,2))
  - ([(2/3(F), 1/3 (C)], [1/3 (F), 2/3 (C)])
    - Expected payoff is (2/3, 2/3)
- The first two are not symmetric; in each one, one of the players yields to the preference of the other player.
- The third equilibrium, in contrast, is symmetric and gives the same payoff to both players, but that payoff is less than 1, the lower payoff in each of the two pure equilibria.

- The players can correlate their actions in the following way. They can toss a fair coin. If the coin comes up heads, they play (F,F), and if it comes up tails they play (C,C). The expected payoff is then (1.5,1.5). Since (F,F) and (C,C) are equilibria, the process we have just described is an equilibrium in an extended game, in which the players can toss a coin and choose their strategies in accordance with the result of the coin toss: after the coin toss, neither player can profit by unilaterally deviating from the strategy recommended by the result of the coin toss.

Correlated equilibrium

- Players' choices of pure strategies may be correlated due to the fact that they use the same random events in deciding which pure strategy to play. Consider an extended game that includes an observer who recommends to each player a pure strategy that he should play. The vector of recommended strategies is chosen by the observer according to a probability distribution over the set of pure strategy vectors, which is commonly known among the players. This probability distribution is called a *correlated equilibrium* if the strategy vector in which all players follow the observer's recommendations is a Nash equilibrium of the extended game.

- The probability distribution over the set of strategy vectors induced by any Nash equilibrium is a correlated equilibrium (though there can be other correlated equilibria too …)
  - Implies directly that correlated equilibrium always exist, since Nash equilibrium exists and each one will correspond to at least one correlated equilibrium.

- The set of correlated equilibria is a polytope that can be calculated as a solution to a set of linear equations.

- Let a denote pure strategy profile, and let a_i denote pure strategy for player i. The variables in the LP are p(a), the probability of realizing a given pure-strategy profile a. Since there is a variable for every pure strategy profile there are thus |A| variables. Observe that as for the two-player zero-sum Nash equilibrium LP, the values $u_i(a)$ are constants.

$$\sum_{a\in A|a_i\in a} p(a)u_i(a) \geq \sum_{a\in A|a_i\in a} p(a)u_i(a'_i, a_{-i}) \quad \forall i \in N, \forall a_i, a'_i \in A_i \quad (4.52)$$

$$p(a) \geq 0 \qquad \forall a \in A \quad (4.53)$$

$$\sum_{a\in A} p(a) = 1 \quad (4.54)$$

Constraints (4.53) and (4.54) ensure that p is a valid probability distribution. The interesting constraint is (4.52), which expresses the requirement that player i must be (weakly) better off playing action a when he is told to do so than playing any other action a'_i, given that other agents play their prescribed actions. This constraint effectively restates the definition of a correlated equilibrium given in Definition 3.4.12. Note that it can be rewritten as $\sum_{a\in A|a_i\in a}[u_i(a) - u_i(a'_i, a_{-i})]p(a) \geq 0$; in other words, whenever agent i is "recommended" to play action a_i with positive probability, he must get at least as much utility from doing so as he would from playing any other action a'_i.

We can select a desired correlated equilibrium by adding an objective function to the linear program. For example, we can find a correlated equilibrium that maximizes the sum of the agents' expected utilities by adding the objective function

$$\text{maximize:} \sum_{a\in A} p(a) \sum_{i\in N} u_i(a). \quad (4.55)$$

Furthermore, all of the questions discussed in Section 4.2.4 can be answered about correlated equilibria in polynomial time, making them (most likely) fundamentally easier problems.

Theorem 4.6.1 *The following problems are in the complexity class P when applied to correlated equilibria:* uniqueness, Pareto optimal, guaranteed payoff, subset inclusion, *and* subset containment.

Finally, it is worthwhile to consider the reason for the computational difference between correlated equilibria and Nash equilibria. Why can we express the definition of a correlated equilibrium as a linear constraint (4.52), while we cannot do the same with the definition of a Nash equilibrium, even though both definitions are quite similar? The difference is that a correlated equilibrium involves a single randomization over action profiles, while in a Nash equilibrium agents randomize separately. Thus, the (nonlinear) version of constraint (4.52) which would instruct a feasibility program to find a Nash equilibrium would be

$$\sum_{a \in A} u_i(a) \prod_{j \in N} p_j(a_j) \geq \sum_{a \in A} u_i(a_i', a_{-i}) \prod_{j \in N \setminus \{i\}} p_j(a_j) \quad \forall i \in N, \forall a_i' \in A_i.$$

This constraint now mimics constraint (4.52), directly expressing the definition of Nash equilibrium. It states that each player i attains at least as much expected utility from following his mixed strategy p_i as from any pure strategy deviation a_i', given the mixed strategies of the other players. However, the constraint is nonlinear because of the product $\prod_{j \in N} p_j(a_j)$.

- One of the underlying assumptions of the concept of equilibrium in strategic-form games is that the choices made by the players are independent. In practice, however, the choices of players may well depend on factors outside the game, and therefore these choices may be correlated. Players can even coordinate their actions among themselves.
  - E.g., in Split or Steal they attempted to correlate their actions in the “negotiation phase.” But was this talk “cheap?”

- One good example of such correlation is the invention of the traffic light: when a motorist arrives at an intersection, he needs to decide whether to cross it, or alternatively to give right of way to motorists approaching the intersection from different directions. If the motorist were to use a mixed strategy in this situation, that would be tantamount to tossing a coin and entering the intersection based on the outcome of the coin toss. If two motorists approaching an intersection simultaneously use this mixed strategy, there is a positive probability that both of them will try to cross the intersection at the same time – which means that there is a positive probability that a traffic accident will ensue. In some states in the US there is an “equilibrium rule” that requires motorists to stop before entering an intersection, and to give right of way to whoever arrived at the intersection earlier. The invention of the traffic light provided a different solution: the traffic light informs each motorist which pure strategy to play, at any given time. The traffic light thus correlates the pure strategies of the players. Note that the traffic light does not, strictly speaking, choose a pure strategy for the motorist; it recommends a pure strategy. It is in the interest of each motorist to follow that recommendation, even if we suppose there are no traffic police watching, no cameras, and no possible court summons awaiting a motorist who disregards the traffic light’s recommendation.

Battle of the sexes

	F	**C**
F	2,1	0,0
C	0, 0	1,2

- The reasoning behind this example is as follows: if we enable the players to conduct a joint (public) lottery, prior to playing the game, they can receive as an equilibrium payoff every convex combination of the equilibrium payoffs of the original game. That is, if we denote by V the set of equilibrium payoffs in the original game, every payoff in the convex hull of V is an equilibrium payoff in the extended game in which the players can conduct a joint lottery prior to playing the game.
- The question naturally arises whether it is possible to create a correlation mechanism, such that the set of equilibrium payoffs in the game that corresponds to this mechanism includes payoffs that are not in the convex hull of V …

Example 8.2 Consider the three-player game depicted in Figure 8.2, in which Player I chooses the row (T or B), Player II chooses the column (L or R), and Player III chooses the matrix (l, c, or r).

l

	L	R
T	0, 1, 3	0, 0, 0
B	1, 1, 1	1, 0, 0

c

	L	R
T	2, 2, 2	0, 0, 0
B	2, 2, 0	2, 2, 2

r

	L	R
T	0, 1, 0	0, 0, 0
B	1, 1, 1	1, 0, 3

Figure 8.2 The payoff matrix of Example 8.2

We will show that the only equilibrium payoff of this game is (1, 1, 1), but there exists a correlation mechanism that induces an equilibrium payoff of (2, 2, 2). In other words, every player gains by using the correlation mechanism. Since (1, 1, 1) is the only equilibrium payoff of the original game, the vector (2, 2, 2) is clearly outside the convex hull of the original game's set of equilibrium payoffs.

Proof sketch

- Step 1: The only equilibrium payoff is (1,1,1).
- See full proof on page 302 of textbook.

- Step 2: The construction of a correlation mechanism leading to the payoff (2,2,2). Consider the following mechanism that the players can implement:
  - Players I and II toss a fair coin, but do not reveal the result of the coin toss to Player III.
  - Players I and II play either (T,L) or (B,R), depending on the result of the coin toss.
  - Player III chooses strategy c.
- Under the implementation of this mechanism, the action vectors that are chosen (with equal probability) are (T,L,c) and (B,R,c), hence the payoff is (2,2,2).
  - Confirm that no player has a unilateral deviation that improves his payoff.

- Note that for the mechanism just described to be an equilibrium, it is necessary that Players I and II know that Player III does not know the result of the coin toss. In other words, while every payoff in the convex hull of the set of equilibrium payoffs can be attained by a *public* lottery, to attain a payoff outside the convex hull of V it is necessary to conduct a lottery that is not public, in which case different players receive different partial information regarding the result of the lottery.

Chicken

	L	R
T	6,6	2,7
B	7, 2	0,0

- The game has three equilibria:
  - (T,R), with payoff (2,7)
  - (B,L), with payoff (7,2)
  - ([2/3(T),1/3(B)],[2/3(L),1/3(R)]), with payoff (4.67,4.67)
- Consider the following mechanism, in which an outside observer gives each player a recommendation regarding which action to take, but the observer does not reveal to either player what recommendation the other player has received. The observer chooses between three action vectors, (T,L), (T,R), (B,L), with equal probability.

	L	**R**
T	1/3	1/3
B	1/3	0

- After conducting a lottery to choose one of the three action vectors, the observer provides Player I with a recommendation to play the first coordinate of that vector. For example, if the action vector (T,L) has been chosen, the observer recommends T to Player I and L to Player II. If Player I receives a recommendation to play T, the conditional probability that Player II has received a recommendation to play L is 1/3 / (1/3 + 1/3) = ½, which is also the conditional probability that he has received a recommendation to play R. In contrast, if Player I receives a recommendation to play B, he knows that Player II has received L as his recommended action.
    - Can show that neither player can profit by a unilateral deviation from the recommendation received from the observer (see page 304 from textbook).
    - Expected equilibrium payoff is (5,5), which lies outside the convex hull of the three equilibrium payoffs of the original game.

- Example shows that the way to attain high payoffs for both players is to avoid the “worst” payoff (0,0). This cannot be accomplished if the players implement independent mixed strategies; it requires correlating the players’ actions. We have made the following assumptions regarding the extended game:
  - The game includes an observer, who recommends strategies.
  - The observer chooses his recommendations probabilistically, based on a distribution commonly known to the players.
  - The recommendations are private, with each player knowing only the recommendation addressed to him or her.
  - The mechanism is common knowledge among the players: each player knows that the mechanism is being used, each player knows that the other players know that the other know that this mechanism is being used, and so forth.

Assignment

- Project proposal (1-2 pages) due on Thursday.
- Reading for next class: chapter 12 from textbook (auctions and mechanism design).