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
- Correlated 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.

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

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

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

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.

Prisoner's dilemma

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

- Suppose row player is the “leader” and column player is the “follower.” What will they play?
- What if column player is “leader” and row player is “follower?”

Stackelberg equilibrium

- A strategy profile (s^S_1, s^S_2) is a *Stackelberg equilibrium* for player 1 if
  - $u_2(s^S_1, s^S_2) >= u_2(s^S_1, s_2)$ for all s_2

  AND
  - $u_1(s^S_1, s^S_2) >= u_1(s_1, s_2)$ for all s_1 and s_2 such that s_2 is a best response to s_1.

- First condition: player 2 is best responding to player 1.
- Second condition: player 1 cannot profitably deviate assuming player 2 will play a best response to his strategy.

- Like extensive-form game where P1 acts first, then P2 acts, and solving for SPE. But we are in strategic-form setting where players act simultaneously.

Chicken

	Swerve	Straight
Swerve	0, 0	-1, +1
Straight	+1, -1	-10, -10

Fig. 2: Chicken with numerical payoffs

Battle of the sexes

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

Rock-paper-scissors

	rock	paper	scissors
Rock	0,0	-1, 1	1, -1
Paper	1,-1	0, 0	-1,1
Scissors	-1,1	1,-1	0,0

Security game

- Random strategy:
  - *Increase cost/uncertainty to attackers*

Adversary

Defender

	Target #1	Target #2
Target #1	4, -3	-1, 1
Target #2	-5, 5	2, -1

	L	R
U	2, 1	4, 0
D	1, 0	3, 1

Commitment to pure strategies

- In this strategic-form representation, the bottom strategy for the row player is strictly dominated by the top strategy. Nevertheless, if the row player has the ability to commit to a pure strategy before the column player chooses his strategy, the row player should commit to the bottom strategy: doing so will make the column player prefer to play the right strategy, leading to a utility of 3 for the row player. By contrast, if the row player were to commit to the top strategy, the column player would prefer to play the left strategy, leading to a utility of only 2 for the row player.

Commitment to mixed strategies

- If the row player commits to placing probability $p > 1/2$ on the bottom strategy, then the column player will still prefer to play the right strategy, and the row player's expected utility will be $3p + 4(1-p) = 4-p \geq 3$.
- If the row player plays each strategy with probability exactly 1/2, the column player is indifferent between the strategies. In such cases, we will assume that the column player will choose the strategy that maximizes the row player's utility (in this case, the right strategy). Hence, the optimal mixed strategy to commit to for the row player is $p=1/2$.

- There are a few good reasons for this assumption. If we were to assume the opposite, then there would not exist an optimal strategy for the row player in the example game: the row player would play the bottom strategy with probability $p=1/2+\varepsilon$ with $\varepsilon>0$,and the smaller ε, the better the utility for the row player. By contrast, if we assume that the follower always breaks ties in the leader's favor, then an optimal mixed strategy for the leader always exists, and this corresponds to a subgame perfect equilibrium of the extensive-form representation of the leadership situation.

- For games with more than two players, in which the players commit to their strategies in sequence, we define optimal strategies to commit to recursively. After the leader commits to a strategy, the game to be played by the remaining agents is itself a (smaller) leadership game. Thus, we define an optimal strategy to commit to as a strategy that maximizes the leader's utility, assuming that the play of the remaining agents is itself optimal under this definition, and maximizes the leader's utility among all optimal ways to play the remaining game. Again, commitment to mixed strategies may or may not be a possibility for every player (although for the last player it does not matter if we allow for commitment to mixed strategies).

Computing the optimal strategy to commit to

- https://www.cs.cmu.edu/~sandholm/Computing%20commitment%20strategy.ec06.pdf

- Theorem: Under commitment to pure strategies, the set of all optimal strategy profiles in a normal-form game can be found in O(#players·#outcomes) time.
- Proof:
  - Each pure strategy that the first player may commit to will induce a subgame for the remaining players. We can solve each such subgame recursively to find all of its optimal strategy profiles; each of these will give the original leader some utility. Those that give the leader maximal utility correspond exactly to the optimal strategy profiles of the original game.

- For general strategic-form games, each player's utility for each of the outcomes has to be explicitly represented in the input, so that the input size is itself $\Omega(\#\text{players} \cdot \#\text{outcomes})$. Therefore, the algorithm is in fact a linear-time algorithm.

Commitment to mixed strategies

- In the special case of two-player zero-sum games, computing an optimal mixed strategy for the leader to commit to is equivalent to computing a minimax strategy, which minimizes the maximum expected utility that the opponent can obtain. Minimax strategies constitute the only natural solution concept for two-player zero-sum games: von Neumann's Minimax Theorem states that in two-player zero-sum games, it does not matter (in terms of the players' utilities) which player gets to commit to a mixed strategy first, and a profile of mixed strategies is a Nash equilibrium if and only if both strategies are minimax strategies. It is well-known that a minimax strategy can be found in polynomial time, using linear programming.

- Theorem: In 2-player strategic-form games, an optimal mixed strategy to commit to can be found in polynomial time using linear programming.

PROOF. For every pure follower strategy t, we compute a mixed strategy for the leader such that 1) playing t is a best response for the follower, and 2) under this constraint, the mixed strategy maximizes the leader's utility. Such a mixed strategy can be computed using the following simple linear program:

$$\begin{array}{l}
\textbf{maximize } \sum_{s \in S} p_s u_l(s,t) \\
\textbf{subject to} \\
\text{for all } t' \in T, \sum_{s \in S} p_s u_f(s,t) \geq \sum_{s \in S} p_s u_f(s,t') \\
\sum_{s \in S} p_s = 1
\end{array}$$

We note that this program may be infeasible for some follower strategies t, for example, if t is a strictly dominated strategy. Nevertheless, the program must be feasible for at least some follower strategies; among these follower strategies, choose a strategy t^* that maximizes the linear program's solution value. Then, if the leader chooses as her mixed strategy the optimal settings of the variables p_s for the linear program for t^*, and the follower plays t^*, this constitutes an optimal strategy profile. □

- Theorem: In 3-player strategic-form games, finding an optimal mixed strategy to commit to is NP-hard.

Security game

	Defender		Attacker	
Target	Covered	Uncovered	Covered	Uncovered
t_1	10	0	-1	1
t_2	0	-10	-1	1

Table 1 Example of a security game with two targets.

- A pure strategy for the defender represents deploying a set of resources on patrols or checkpoints, e.g., scheduling checkpoints at the LAX airport or assigning federal air marshals to protect flight tours. The pure strategy for an attacker represents an attack at a target, e.g., a flight. The strategy for the leader is a mixed strategy, a probability distribution over the pure strategies of the defender. Additionally, with each target are also associated a set of payoff values that define the utilities for both the defender and the attacker in case of a successful or a failed attack.

- A key assumption of security games is that the payoff of an outcome depends only on the target attacked, and whether or not it is covered by the defender. The payoffs do not depend on the remaining aspects of the defender allocation. For example, if an adversary succeeds in attacking target t_1, the penalty for the defender is the same whether the defender was guarding target t_2, or not.
- This allows us to compactly represent the payoffs of a security game. Specifically, a set of four payoffs is associated with each target. These four payoffs are the rewards and penalties to both the defender and the attacker in case of a successful or an unsuccessful attack, and are sufficient to define the utilities for both players for all possible outcomes in the security domain.

Security game

	Defender		Attacker	
Target	Covered	Uncovered	Covered	Uncovered
t_1	10	0	-1	1
t_2	0	-10	-1	1

Table 1 Example of a security game with two targets.

Strong Stackelberg equilibrium

- A pair of strategies form a *Strong Stackelberg Equilibrium (SSE)* if they satisfy
  1. The defender plays a best response. That is, the defender cannot get a higher payoff by choosing any other strategy.
  2. The attacker plays a best response. That is, given a defender strategy, the attacker cannot get a higher payoff by attacking any other target.
  3. The attacker breaks ties in favor of the leader.

- The assumption that the follower will always break ties in favor of the leader in cases of indifference is reasonable because in most cases the leader can induce the favorable strong equilibrium by selecting a strategy arbitrarily close to the equilibrium that causes the follower to strictly prefer the desired strategy. Furthermore an SSE exists in all Stackelberg games, which makes it an attractive solution concept compared to versions of Stackelberg equilibrium with other tie-breaking rules.

Applications of security games

- ARMOR for Los Angeles International airport
  - http://teamcore.usc.edu/, http://create.usc.edu/sites/default/files/publications/computationalgametheoryforsecurityandsustainability.pdf
  - Vehicular checkpoints, police units patrolling the roads to the terminals, patrolling inside the terminals (with canines), and security screening and bag checks for passengers.
  - Need to allocate resources to eight different terminals with very different characteristics (physical size, passenger load, international vs. domestic flights, etc.)
  - Assume n roads. Police strategy to place $m < n$ checkpoints on the roads, where m is maximum number of checkpoints.
  - ARMOR randomizes allocation of checkpoints to roads.
  - The adversary may conduct surveillance of this mixed strategy and potentially choose to act through one of the roads.

IRIS for US Federal Air Marshall Service

- The US Federal Air Marshals Service allocates air marshals to flights originating in and departing from the United States to dissuade potential aggressors and prevent an attack should one occur. Flights are of different importance based on a variety of factors such as the numbers of passengers, the population of source and destination, and international flights from different countries. Security resource allocation in this domain is significantly more challenging than for ARMOR: a limited number of air marshals need to be scheduled to cover thousands of commercial flights each day. Furthermore, these air marshals must be scheduled on tours of flights that obey various constraints (e.g., the time required to board, fly, and disembark). Simply finding schedules for the marshals that meet all of these constraints is a computational challenge. Our task is made more difficult by the need to find a randomized policy that meets these scheduling constraints, while also accounting for the different values of each flight.

- PROTECT for US Coast Guard
  - Given a particular port and the variety of critical infrastructure that an attacker may attack within the port, USCG conducts patrols to protect this infrastructure; however, while the attacker has the opportunity to observe patrol patterns, limited security resources imply that USCG patrols cannot be at every location 24/7.
  - It has been in use at the port of Boston since April 2011, and is also in use at the port of New York since February 2012. Similar to previous applications ARMOR and IRIS, PROTECT uses an attacker-defender Stackelberg game framework, with USCG as the defender against terrorists that conduct surveillance before potentially launching an attack.
- Ferry Protection for the US Coast Guard
- TRUSTS for Security in Transit Systems

Assignment

- Project proposal (1-2 pages) due on Tuesday.
- Reading for next class: chapter 10 from Bauso textbook.