CAP 5993/CAP 4993
Game Theory

Instructor: Sam Ganzfried

sganzfri@cis.fiu.edu

Schedule

- HW4 out this week due 4/13.
- 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

- 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
- Stackelberg equilibrium

Game representations

- Strategic form
- Extensive form
  - Perfect information
  - Perfect information (with chance events)
  - Imperfect information (with chance events)
- Repeated
  - Finite vs. infinite
  - Discounted vs. undiscounted

The "Big Match"

- One day, the king has to leave for an undefined time and therefore decides to put his trusted minister in charge of the kingdom. The day before leaving, the king informs the minister that he will not hear from the king until his return. On the day the king will return, if the minister will be found *working hard*, the king will award the minister by abducting in favor of him. On the other hand, if on that day the king will find the minister *enjoying life*, the king will put the minister in prison. The king is powerful and has informers. Therefore he knows every day whether the minister was at work or not in the past days.

- The minister knows that if he worked hard every day, the king, being informed of this, would not come back. But this would mean an everlasting miserable life of working hard every day!

- The minister also knows that if he did not work hard at all, the king would come very soon and the minister would be imprisoned.

	L	R
T	0 s_2	1 s_2
B	1 s_1	0 s_0

State s_2

	L
T	1 s_1

State s_1

	L
T	0 s_0

State s_0

- The row player is the king, and the column player is the minister. The decision of the king not to come back corresponds to action T. Thus, for every day that the king plays T, the state of the game transitions to the same state s_2. This occurs independently of the choice of the minister to be at work, denote by L, or to rest, denoted by R. The choice of the king of coming back is denote by the action B. If the king plays B and the minster plays L (work hard), the game jumps to state s_1, which implies an everlasting reward for the minister. Conversely, if the king plays B and the minister plays R (the minister is found enjoying life), then the state of the game jumps to s_0, which implies an everlasting punishment for the minister.

- We can solve this game by applying “dynamic programming principle.” Let λ be the discount factor. Suppose v_i are the “values” of the stage games. Then we can compute the values of each action in each game as a function of v_i as follows:

	L	R
T	$(1-\lambda)v_2$	$\lambda+(1-\lambda)v_2$
B	$\lambda+(1-\lambda)v_1$	$(1-\lambda)v_0$

Game $G_{s_2}^{\lambda}$

	L
T	$\lambda+(1-\lambda)v_1$

Game $G_{s_1}^{\lambda}$

	L
T	$(1-\lambda)v_0$

Game $G_{s_0}^{\lambda}$

- By imposing the fixed point condition on both states 0 and 1, we obtain: $v_0 = 0$, $v_1 = 1$.
- From the Indifference Principle, we can then solve for v_2 to obtain $v_2 = ½$.
- The equilibrium strategies are [½ L, ½ R] for row player and [1/(1+ λ) T, λ /(1+ λ) B] for column player.
- (Full derivation in Bauso textbook).

- The interpretation of the above result is as follows. The best strategy for the minister is to work every two days on average. This is equivalent to saying that every day the minister will toss a coin and depending on the result he will work hard or not. The interpretation of the best strategy for the king is as follows. First note that his optimal strategy will depend on the discount factor, that is, on how farsighted he is. The king will return with a probability that increases with the discount factor. That is to say that the more myopic the king is, the sooner he will come back. Conversely, if the king is farsighted, the discount factor is small and tends to zero, and consequently the probability of coming back approaches zero. Note that the discount factor influences only the strategy of the king. The strategy of the minister does not depend on the discount factor. This derives from the fact that only the king can force the state of the game to jump to an absorbing state.

Stochastic games

Definition 6.2.1 (Stochastic game) *A* stochastic game *(also known as a* Marko[v] game*) is a tuple* (Q, N, A, P, r), *where:*

- Q *is a finite set of games;*
- N *is a finite set of* n *players;*
- $A = A_1 \times \cdots \times A_n$, *where* A_i *is a finite set of actions available to player* i;
- $P : Q \times A \times Q \mapsto [0, 1]$ *is the transition probability function;* $P(q, a, \hat{q})$ *is the probability of transitioning from state* q *to state* $\hat{q}$ *after action profile* a; *and*
- $R = r_1, \ldots, r_n$, *where* $r_i : Q \times A \mapsto \mathbb{R}$ *is a real-valued payoff function for player* i.

In this definition we have assumed that the strategy space of the agents is the same in all games, and thus that the difference between the games is only in the payoff function. Removing this assumption adds notation, but otherwise presents no major difficulty or insights. Restricting Q and each A_i to be finite is a substantive restriction, but we do so for a reason; the infinite case raises a number of complications that we wish to avoid.

We have specified the payoff of a player at each stage game (or in each state), but not how these payoffs are aggregated into an overall payoff. To solve this problem, we can use solutions already discussed earlier in connection with infinitely repeated games (Section 6.1.2). Specifically, the two most commonly used aggregation methods are *average reward* and *future discounted reward*.

Stochastic games generalize many settings

- Games with *finite interactions*; this occurs if the state of the game reaches at time t an absorbing state with null payoff;
- Static *matrix games* (aka strategic-form games) if we set t = 1
- *Repeated games* if the game admits only one state
- *Stopping games* if the stage payoff is null until a player decides to *quit* the game; in consequence of this, the state of the game reaches an absorbing state with normal payoff.
- *Markov decision problems* if the game involves only one single player.
  - We will see an example later.

Stochastic (aka Markov) games

- Capital accumulation or fishery:
- Taxation:
- Communication network:
- Queues:
- Poker tournament:
  - Stacks of (20,50,10) -> (30,40,10) ->…

- Theorem (Shapley 1953): If all sets are finite, then for every λ there exists an equilibrium in stationary strategies.
  - Proof: Uses the above "dynamic programming" procedure, where "nonexpansiveness" of the value operator yields a unique fixed point, which corresponds to a Nash equilibrium.
- A strategy is *stationary* if it depends only on the current state (and not on the time step).

- Theorem (Mertens and Neyman 1981): For two-player zero-sum games, each player has a strategy that is ε-optimal for every discount factor sufficiently small.
  - Called a "uniform equilibrium"

- Theorem (Vielle 2000): For every two-player nonzero-sum stochastic game there is a strategy profile that is an ε-equilibrium for every discount factor sufficiently small.

Continuous games

- $G = (P,C,U)$
  - $P = 1,2,3,..,n$ is the set of players
  - $C = (C_1,\ldots,C_n)$ is a compact metric space corresponding to the i'th player's set of pure strategies
  - $U = (u_1,\ldots,u_n)$ is utility function of player i

- Mixed strategies are the space of Borel probability measures on C_i.
- The existence of a Nash equilibrium for any continuous game with continuous utility functions can been proven using Irving Glicksberg's generalization of the Kakutani fixed point theorem. In general, there may not be a solution if we allow strategy spaces, C_i's which are not compact, or if we allow non-continuous utility functions.

Continuous games

- It extends the notion of a discrete game, where the players choose from a finite set of pure strategies. The continuous game concepts allows games to include more general sets of pure strategies, which may be uncountably infinite.
- In general, a game with uncountably infinite strategy sets will not necessarily have a Nash equilibrium solution. If, however, the strategy sets are required to be compact and the utility functions continuous, then a Nash equilibrium will be guaranteed; this is by Glicksberg's generalization of the Kakutani fixed point theorem. The class of continuous games is for this reason usually defined and studied as a subset of the larger class of infinite games (i.e. games with infinite strategy sets) in which the strategy sets are compact and the utility functions continuous.

Guess the larger number

- What are the strategy sets and utility functions?
- Why does this have no equilibrium??

Separable continuous games

Separable games [edit]

A **separable game** is a continuous game where, for any i, the utility function $u_i : \mathbf{C} \to \mathbb{R}$ can be expressed in the sum-of-products form:

$$u_i(\mathbf{s}) = \sum_{k_1=1}^{m_1} \ldots \sum_{k_n=1}^{m_n} a_{i\,,\,k_1\ldots k_n} f_1(s_1)\ldots f_n(s_n)$$, where $\mathbf{s} \in \mathbf{C}$, $s_i \in C_i$, $a_{i\,,\,k_1\ldots k_n} \in \mathbb{R}$, and the functions $f_{i\,,\,k} : C_i \to \mathbb{R}$ are continuous.

A **polynomial game** is a separable game where each C_i is a compact interval on $\mathbb{R}$ and each utility function can be written as a multivariate polynomial.

In general, mixed Nash equilibria of separable games are easier to compute than non-separable games as implied by the following theorem:

> For any separable game there exists at least one Nash equilibrium where player *i* mixes at most $m_i + 1$ pure strategies.[2]

Whereas an equilibrium strategy for a non-separable game may require an uncountably infinite support, a separable game is guaranteed to have at least one Nash equilibrium with finitely supported mixed strategies.

Separable games [edit]

A polynomial game [edit]

Consider a zero-sum 2-player game between players **X** and **Y**, with $C_X = C_Y = [0,1]$. Denote elements of C_X and C_Y as x and y respectively. Define the utility functions $H(x,y) = u_x(x,y) = -u_y(x,y)$ where

$$H(x,y) = (x-y)^2 .$$

The pure strategy best response relations are:

$$b_X(y) = \begin{cases} 1, & \text{if } y \in [0,1/2) \\ 0 \text{ or } 1, & \text{if } y = 1/2 \\ 0, & \text{if } y \in (1/2,1] \end{cases}$$

$$b_Y(x) = x$$

$b_X(y)$ and $b_Y(x)$ do not intersect, so there is

no pure strategy Nash equilibrium. However, there should be a mixed strategy equilibrium. To find it, express the expected value, $v = \mathbb{E}[H(x,y)]$ as a linear combination of the first and second moments of the probability distributions of **X** and **Y**:

$$v = \mu_{X2} - 2\mu_{X1}\mu_{Y1} + \mu_{Y2}$$

(where $\mu_{XN} = \mathbb{E}[x^N]$ and similarly for *Y*).

The constraints on μ_{X1} and μ_{X2} (with similar constraints for *y*,) are given by Hausdorff as:

$$\begin{aligned} \mu_{X1} &\geq \mu_{X2} \qquad & \mu_{Y1} &\geq \mu_{Y2} \\ \mu_{X1}^2 &\leq \mu_{X2} \qquad & \mu_{Y1}^2 &\leq \mu_{Y2} \end{aligned}$$

Each pair of constraints defines a compact convex subset in the plane. Since v is linear, any extrema with respect to a player's first two moments will lie on the boundary of this subset. Player i's equilibrium strategy will lie on

$$\mu_{i1} = \mu_{i2} \text{ or } \mu_{i1}^2 = \mu_{i2}$$

Note that the first equation only permits mixtures of 0 and 1 whereas the second equation only permits pure strategies. Moreover, if the best response at a certain point to player i lies on $\mu_{i1} = \mu_{i2}$, it will lie on the whole line, so that both 0 and 1 are a best response. $b_Y(\mu_{X1}, \mu_{X2})$ simply gives the pure strategy $y = \mu_{X1}$, so b_Y will never give both 0 and 1. However b_x gives both 0 and 1 when y = 1/2. A Nash equilibrium exists when:

$$(\mu_{X1}*, \mu_{X2}*, \mu_{Y1}*, \mu_{Y2}*) = (1/2, 1/2, 1/2, 1/4)$$

This determines one unique equilibrium where Player X plays a random mixture of 0 for 1/2 of the time and 1 the other 1/2 of the time. Player Y plays the pure strategy of 1/2. The value of the game is 1/4.

Non-Separable Games [edit]

A rational pay-off function [edit]

Consider a zero-sum 2-player game between players **X** and **Y**, with $C_X = C_Y = [0,1]$. Denote elements of C_X and C_Y as x and y respectively. Define the utility functions $H(x,y) = u_x(x,y) = -u_y(x,y)$ where

$$H(x,y) = \frac{(1+x)(1+y)(1-xy)}{(1+xy)^2}.$$

This game has no pure strategy Nash equilibrium. It can be shown[3] that a unique mixed strategy Nash equilibrium exists with the following pair of probability density functions:

$$f^*(x) = \frac{2}{\pi\sqrt{x}(1+x)} \qquad g^*(y) = \frac{2}{\pi\sqrt{y}(1+y)}.$$

The value of the game is $4/\pi$.

Requiring a Cantor distribution [edit]

Consider a zero-sum 2-player game between players **X** and **Y**, with $C_X = C_Y = [0,1]$. Denote elements of C_X and C_Y as x and y respectively. Define the utility functions $H(x,y) = u_x(x,y) = -u_y(x,y)$ where

$$H(x,y) = \sum_{n=0}^{\infty} \frac{1}{2^n} \left(2x^n - \left(\left(1-\frac{x}{3}\right)^n - \left(\frac{x}{3}\right)^n\right)\right) \left(2y^n - \left(\left(1-\frac{y}{3}\right)^n - \left(\frac{y}{3}\right)^n\right)\right).$$

This game has a unique mixed strategy equilibrium where each player plays a mixed strategy with the cantor singular function as the cumulative distribution function.[4]

Continuous poker models

- Consider the following simplified poker game. Suppose two players are given private signals, x_1 and x_2, independently and uniformly at random from [0,1]. Suppose the pot initially has size p (one can think of this as both players having put in an ante of p/2, or that we are at the final betting round—aka final street—of a multi-street game). Player 1 is allowed to bet or check. If player 1 checks, the game is over and the player with the lower private signal wins the pot (following the convention of). If player 1 bets, then player 2 can call or fold. If player 2 folds, then player 1 wins the pot. If player 2 calls, then whoever has the lower private signal wins p+1, while the other player loses 1. This situation can be thought of as an abstraction of the final street of a hand of limit Texas hold ’em where raising is not allowed and player 2 has already checked.

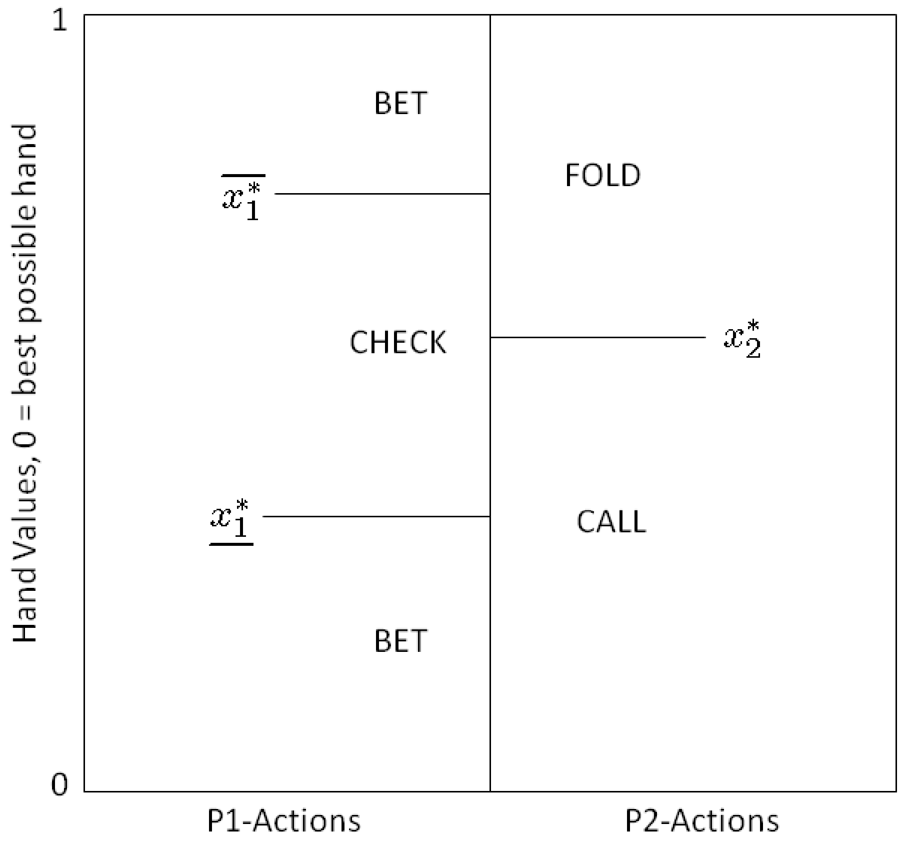
1
0
Hand Values, 0 = best possible hand
BET
$\overline{x_1^*}$
CHECK
$\underline{x_1^*}$
BET
FOLD
x_2^*
CALL
P1-Actions
P2-Actions

- These poker game models are generally not separable, but can often be solved analytically by solving a series of indifference equations.
  - Many examples of this in “Mathematics of Poker” by Ankenman and Chen.
  - See also my paper https://www.cs.cmu.edu/~sganzfri/Qualitative_AAMAS10.pdf (extended version here https://www.cs.cmu.edu/~sganzfri/Qualitative_TR10.pdf).

Bayesian games

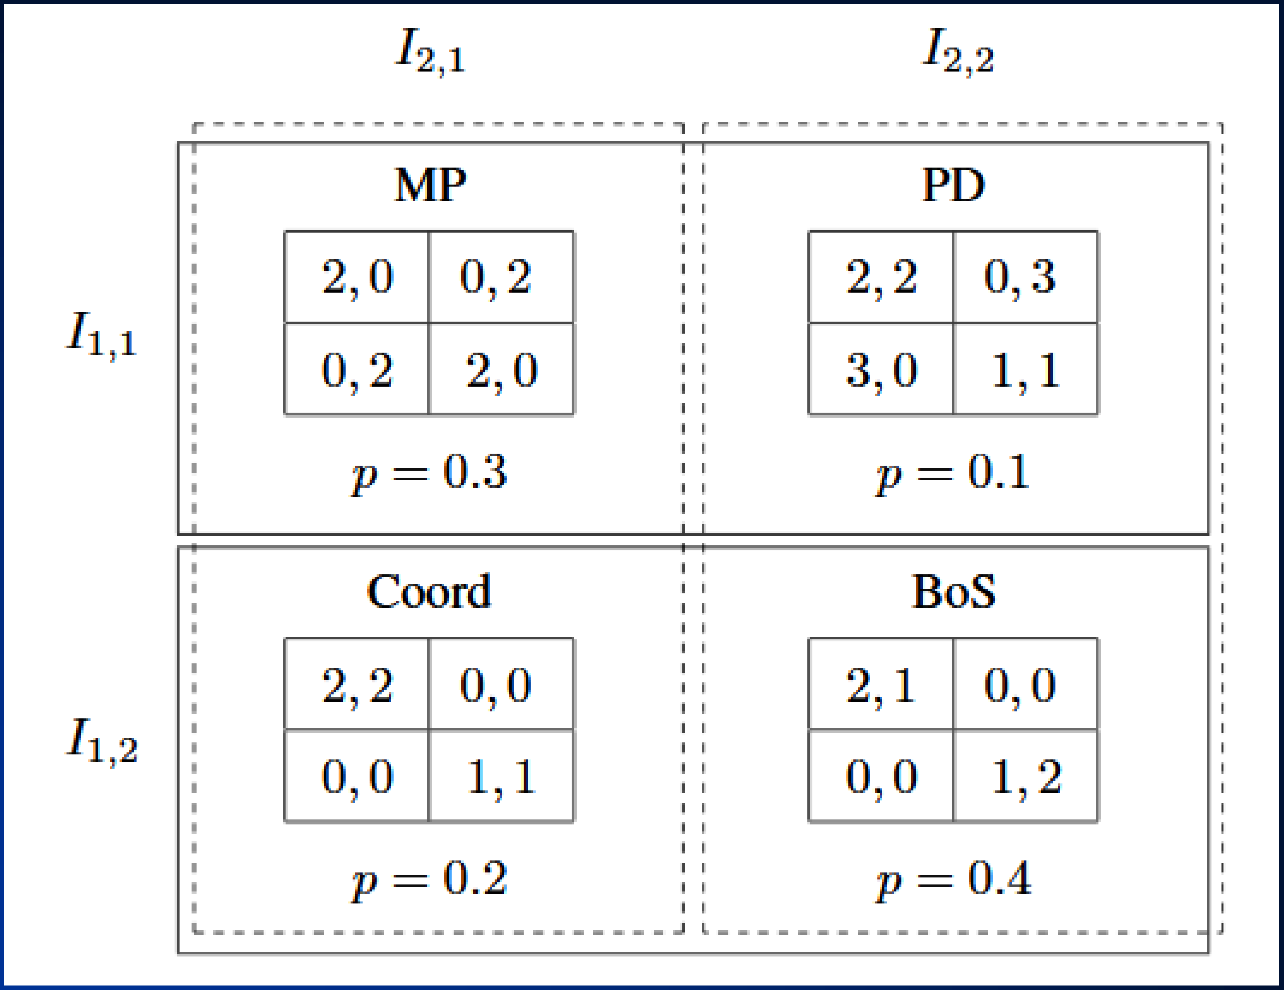

- Players have uncertainty about the very game being played!!!
  - "incomplete information"
- Contrast this with when players knew the game but not the exact state of the game, "imperfect information"

- Two assumptions:
  - All possible games have the same number of agents and the same strategy space for each agent; they differ only in their payoffs
  - (common prior assumption) The beliefs of the different agents are posteriors, obtained by conditioning a common prior on individual prior signals

3 (equivalent) definitions

Definition 6.3.1 (Bayesian game: information sets) *A* Bayesian game *is a tuple (N, G, P, I) where:*

- *N is a set of agents;*
- *G is a set of games with N agents each such that if $g, g' \in G$ then for each agent $i \in N$ the strategy space in g is identical to the strategy space in g';*
- *$P \in \Pi(G)$ is a common prior over games, where $\Pi(G)$ is the set of all probability distributions over G; and*
- *$I = (I_1, ..., I_N)$ is a tuple of partitions of G, one for each agent.*

choice be interspersed arbitrarily with the agents' moves, without loss of generality we assume that Nature makes all its choices at the outset. Nature does not have a utility function (or, alternatively, can be viewed as having a constant one), and has the unique strategy of randomizing in a commonly known way. The agents receive individual signals about Nature's choice, and these are captured by their information sets in a standard way. The agents have no additional information; in particular, the information sets capture the fact that agents make their choices without knowing the choices of others. Thus, we have reduced games of incomplete information to games of imperfect information, albeit ones with chance moves. These chance moves of Nature require minor adjustments of existing definitions, replacing payoffs by their expectations given Nature's moves.[6]

For example, the Bayesian game of Figure 6.7 can be represented in extensive form as depicted in Figure 6.8.

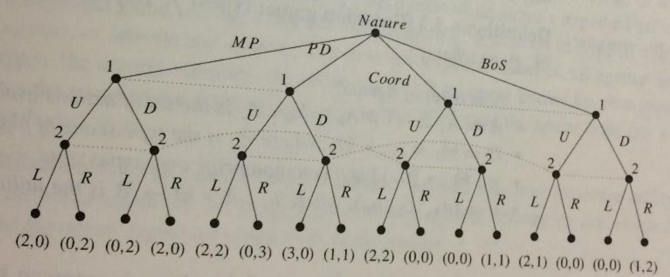

Figure 6.8 The Bayesian game from Figure 6.7 in extensive form.

Although this second definition of Bayesian games can be initially more intuitive than our first definition, it can also be more cumbersome to work with. This is because we use an extensive-form representation in a setting where players are unable to observe each others' moves. (Indeed, for the same reason we do not routinely use extensive-form games of imperfect information to model simultaneous interactions such as the Prisoner's Dilemma, though we could do so if we wished.) For this reason, we will not make further use of this definition. We close by noting one advantage that it does have, however: it extends very naturally ... players move sequentially and do (at least

a_1	a_2	θ_1	θ_2	u_1	u_2
U	L	$\theta_{1,1}$	$\theta_{2,1}$	2	0
U	L	$\theta_{1,1}$	$\theta_{2,2}$	2	2
U	L	$\theta_{1,2}$	$\theta_{2,1}$	2	2
U	L	$\theta_{1,2}$	$\theta_{2,2}$	2	1
U	R	$\theta_{1,1}$	$\theta_{2,1}$	0	2
U	R	$\theta_{1,1}$	$\theta_{2,2}$	0	3
U	R	$\theta_{1,2}$	$\theta_{2,1}$	0	0
U	R	$\theta_{1,2}$	$\theta_{2,2}$	0	0

a_1	a_2	θ_1	θ_2	u_1	u_2
D	L	$\theta_{1,1}$	$\theta_{2,1}$	0	2
D	L	$\theta_{1,1}$	$\theta_{2,2}$	3	0
D	L	$\theta_{1,2}$	$\theta_{2,1}$	0	0
D	L	$\theta_{1,2}$	$\theta_{2,2}$	0	0
D	R	$\theta_{1,1}$	$\theta_{2,1}$	2	0
D	R	$\theta_{1,1}$	$\theta_{2,2}$	1	1
D	R	$\theta_{1,2}$	$\theta_{2,1}$	1	1
D	R	$\theta_{1,2}$	$\theta_{2,2}$	1	2

Figure 6.9 Utility functions u_1 and u_2 for the Bayesian game from Figure 6.7.

Our third definition uses the notion of an *epistemic type*, or simply a *type* as a way of defining uncertainty directly over a game's utility function.

Definition 6.3.2 (Bayesian game: types) *A* Bayesian game *is a tuple* (N, A, Θ, p, u) *where:*

- *N is a set of agents;*
- *$A = A_1 \times \cdots \times A_n$, where A_i is the set of actions available to player i;*
- *$\Theta = \Theta_1 \times \ldots \times \Theta_n$, where Θ_i is the type space of player i;*
- *$p : \Theta \mapsto [0, 1]$ is a common prior over types; and*
- *$u = (u_1, \ldots, u_n)$, where $u_i : A \times \Theta \mapsto \mathbb{R}$ is the utility function for player i.*

The assumption is that all of the above is common knowledge among the players, and that each agent knows his own type. This definition can seem mysterious, because the notion of type can be rather opaque. In general, the type of agent encapsulates all the information possessed by the agent that is not common knowledge. This is often quite simple (e.g., the agent's knowledge of his private payoff function), but can also include his beliefs about other agents' payoffs, about their beliefs about his own payoff, and any other higher-order beliefs.

We can get further insight into the notion of a type by relating it to the formulation at the beginning of this section. Consider again the Bayesian game in Figure 6.7. For each of the agents we have two types, corresponding to his two information sets. Denote player 1's actions as U and D, player 2's actions as L and R. Call the types of the first agent $\theta_{1,1}$ and $\theta_{1,2}$, and those of the second agent $\theta_{2,1}$ and $\theta_{2,2}$. The joint distribution on these types is as follows: $p(\theta_{1,1}, \theta_{2,1}) = .3$, $p(\theta_{1,1}, \theta_{2,2}) = .1$, $p(\theta_{1,2}, \theta_{2,1}) = .2$, $p(\theta_{1,2}, \theta_{2,2}) = .4$. The conditional probabilities for the first player are $p(\theta_{2,1} \mid \theta_{1,1}) = 3/4$, $p(\theta_{2,2} \mid \theta_{1,1}) = 1/4$, $p(\theta_{2,1} \mid \theta_{1,2}) = 1/3$, and $p(\theta_{2,2} \mid \theta_{1,2}) = 2/3$. Both players' utility functions are given in Figure 6.9.

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.

Bayesian Security games

- In the above example, all payoff values are exactly known. In practice, we often have uncertainty over the payoffs and preferences of the players. Bayesian games are a well-known game-theoretic model in which such uncertainty is modeled using multiple types of players, with each associated with its own payoff values. For security games of interest, the main source of payoff uncertainty is regarding the attacker's payoffs. In the resulting Bayesian Stackelberg game model, there is only one leader type (e.g., only one police force), although there can be multiple follower types (e.g., multiple attacker types trying to infiltrate security). Each follower type is represented using a different payoff matrix. The leader does not know the follower's type, but knows the probability distribution over them. The goal is to find the optimal mixed strategy for the leader to commit to, given that the defender could be facing any of the follower types.

Robust management of diabetes

Assignment

- Project proposal (1-2 pages) due tonight.
- Reading for next class: chapter 12 from main textbook.
- Homework 4 out this week (due 4/13).