Cap. 2. Introdução a Teoria de Utilidade e Teoria de Jogos

Teoria de Utilidade

- Como as pessoas escolhem e como formalizar o processo?
  - Preferências Racionais
  - Utilidade
  - Utilidade x Dinheiro
  - Utilidades multiatributos

Lotteries

An agent chooses among prizes (A, B, etc.) and lotteries, i.e., situations with uncertain prizes

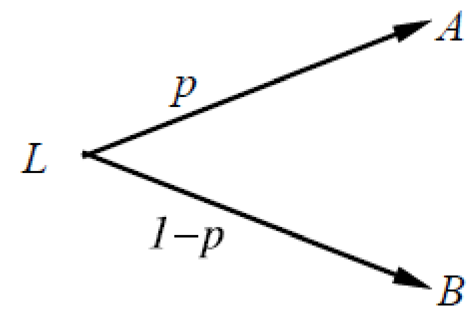

Lottery $L = [p, A;\ (1 - p), B]$

Notation:

$A \succ B$	A preferred to B
$A \sim B$	indifference between A and B
$A \succsim B$	B not preferred to A

Preferências Racionais

Idea: preferences of a rational agent must obey constraints.
Rational preferences $\Rightarrow$
behavior describable as maximization of expected utility

Constraints:

Orderability

$$(A \succ B) \vee (B \succ A) \vee (A \sim B)$$

Transitivity

$$(A \succ B) \wedge (B \succ C) \Rightarrow (A \succ C)$$

Continuity

$$A \succ B \succ C \Rightarrow \exists p \ [p, A; \ 1-p, C] \sim B$$

Substitutability

$$A \sim B \Rightarrow [p, A; \ 1-p, C] \sim [p, B; 1-p, C]$$

Monotonicity

$$A \succ B \Rightarrow ([p, A; \ 1-p, B] \succsim [q, A; \ 1-q, B] \Leftrightarrow \ p \geq q)$$

Violação das Restrições leva a Irracionalidade

For example: an agent with intransitive preferences can be induced to give away all its money

If $B \succ C$, then an agent who has C would pay (say) 1 cent to get B

If $A \succ B$, then an agent who has B would pay (say) 1 cent to get A

If $C \succ A$, then an agent who has A would pay (say) 1 cent to get C

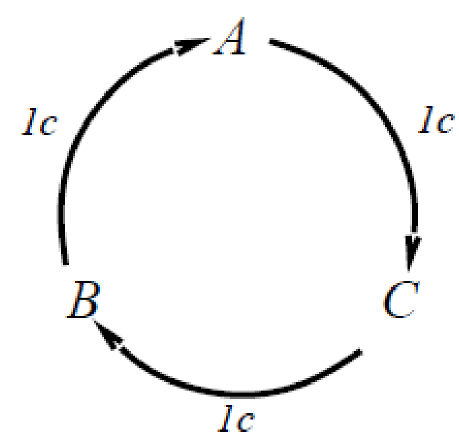

Maximizing Expected Utility (MEU)

Theorem (Ramsey, 1931; von Neumann and Morgenstern, 1944):
Given preferences satisfying the constraints
there exists a real-valued function U such that

$$U(A) \geq U(B) \quad \Leftrightarrow \quad A \succsim B$$
$$U([p_1, S_1; \; \ldots \; ; \; p_n, S_n]) = \Sigma_i \; p_i U(S_i)$$

MEU principle:
Choose the action that maximizes expected utility

Note: an agent can be entirely rational (consistent with MEU) without ever representing or manipulating utilities and probabilities

E.g., a lookup table for perfect tictactoe

Utilidades

Utilities map states to real numbers. Which numbers?

Standard approach to assessment of human utilities:

compare a given state A to a standard lottery L_p that has

"best possible prize" $u_\top$ with probability p

"worst possible catastrophe" $u_\perp$ with probability $(1-p)$

adjust lottery probability p until $A \sim L_p$

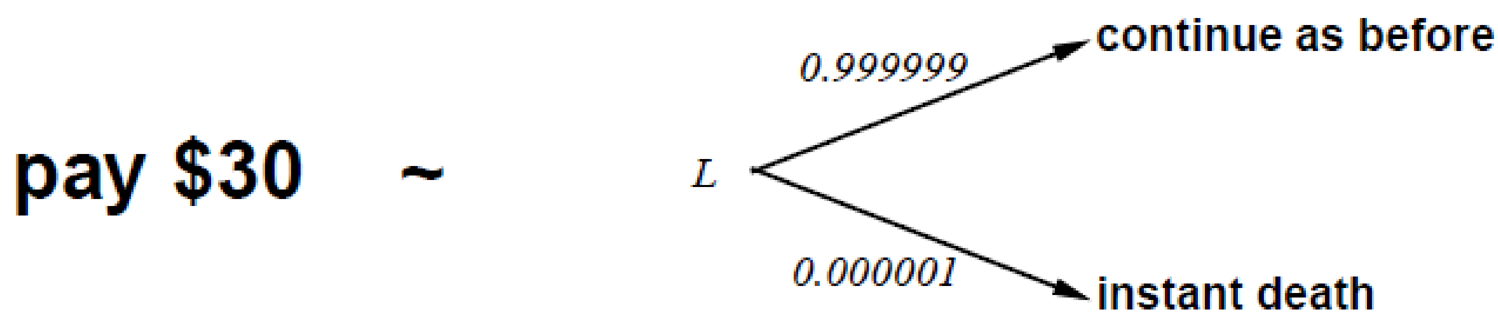

Definindo Funções de Utilidades através de loteria

Normalized utilities: $u_\top = 1.0$, $u_\bot = 0.0$

Micromorts: one-millionth chance of death
useful for Russian roulette, paying to reduce product risks, etc.

QALYs: quality-adjusted life years
useful for medical decisions involving substantial risk

Note: behavior is invariant w.r.t. +ve linear transformation

$$U'(x) = k_1 U(x) + k_2 \quad \text{where } k_1 > 0$$

With deterministic prizes only (no lottery choices), only ordinal utility can be determined, i.e., total order on prizes

- Funções de utilidades ordinais podem ser chamadas de funções de valor e são invariantes para qualquer transformação monotônica

Preferências de indivíduos sobre dinheiro certo (x) e loteria [p,M;1-p,0]

For each x, adjust p until half the class votes for lottery (M=10,000)

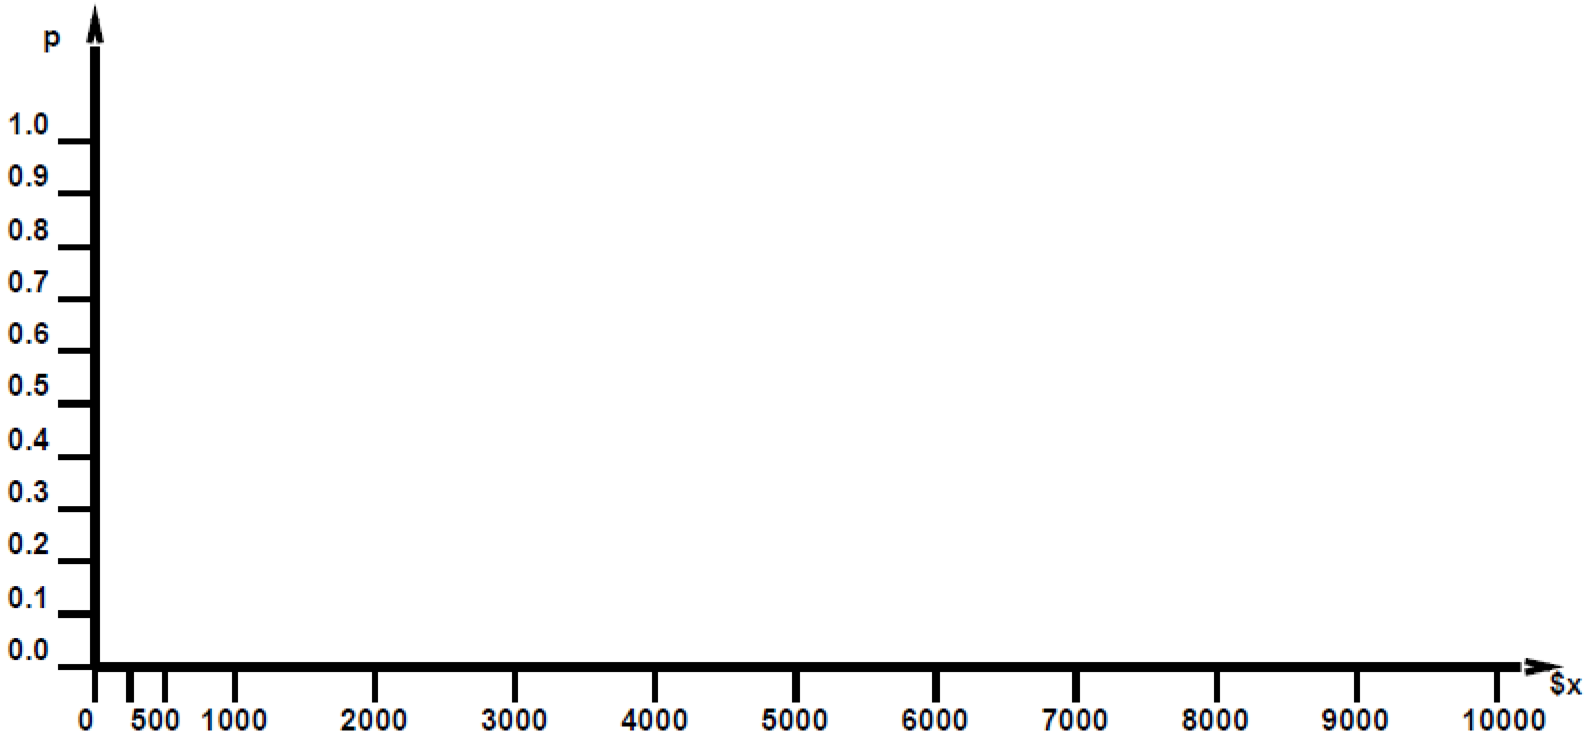

Money vs Utility

Money does not behave as a utility function

Given a lottery L with expected monetary value $EMV(L)$, usually $U(L) < U(EMV(L))$, i.e., people are risk-averse

Utility curve: for what probability p am I indifferent between a fixed prize x and a lottery $[p, \$M; (1-p), \$0]$ for large M?

Typical empirical data, extrapolated with risk-prone behavior:

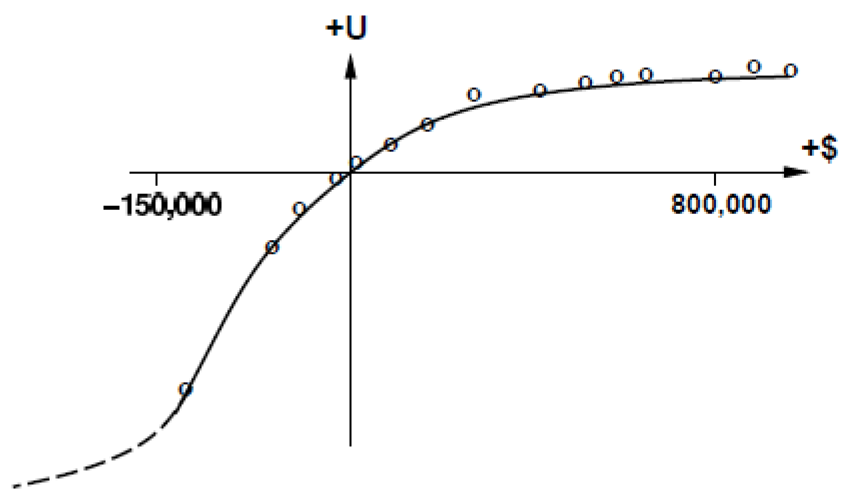

The Saint Petersburg Paradox

- The paradox is named from Daniel Bernoulli's presentation of the problem and his solution, published in 1738 in St. Petersburg

- A casino offers a game of chance for a single player in which a fair coin is tossed at each stage. The pot starts at 1 dollar and is doubled every time a head appears. The first time a tail appears, the game ends and the player wins whatever is in the pot.

- Thus the player wins 1 dollar if a tail appears on the first toss, 2 dollars if a head appears on the first toss and a tail on the second.
- Two questions:
  - How much would you accept to pay for playing this game?
  - What is the expected monetary value of the game?

The Saint Petersburg Paradox

- As Bernoulli stated:
  - The determination of the value of an item must not be based on the price, but rather on the utility it yields…. There is no doubt that a gain of one thousand ducats is more significant to the pauper than to a rich man though both gain the same amount
- Bernoulli proposed that utility of money should be logarithmic. U(M)= a*log2(M)+b
- This makes EMV to be a finite value.

- But it's always possible to recreate the paradox by changing the function!!!
  - Alternative theories may provide a better description model (*Prospect Theory*)

Multiatibute utility functions

How can we handle utility functions of many variables $X_1 \ldots X_n$?
E.g., what is $U(Deaths, Noise, Cost)$?

How can complex utility functions be assessed from preference behaviour?

Idea 1: identify conditions under which decisions can be made without complete identification of $U(x_1, \ldots, x_n)$

Idea 2: identify various types of independence in preferences and derive consequent canonical forms for $U(x_1, \ldots, x_n)$

Airport site example

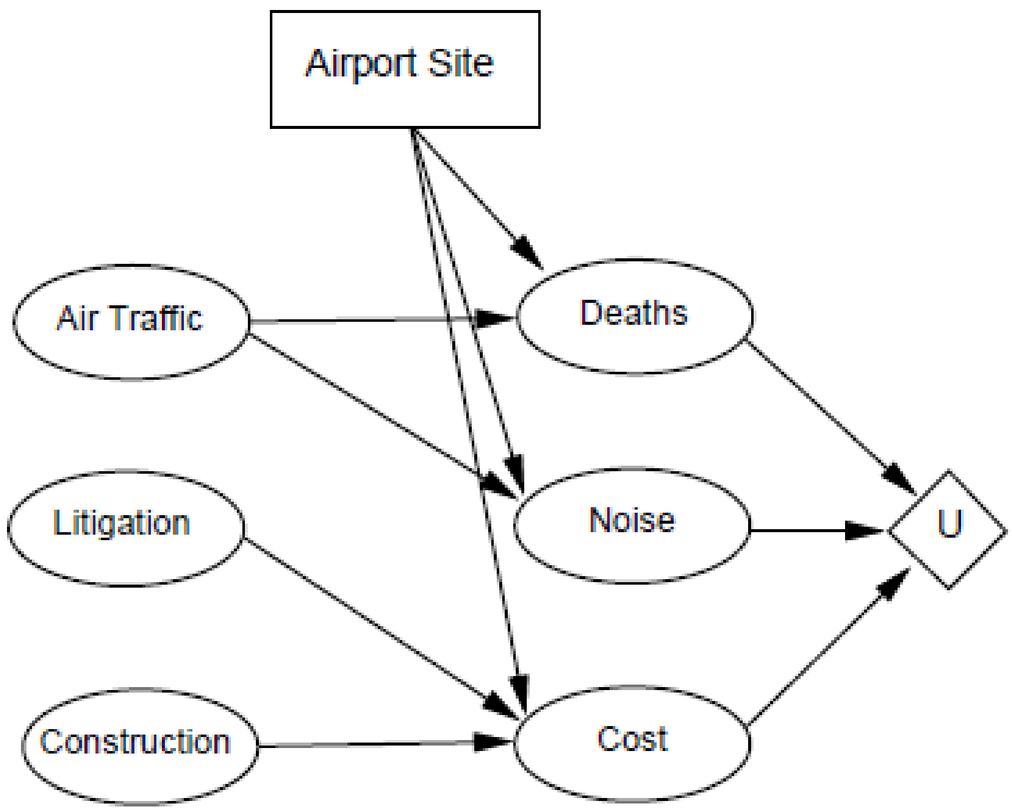

Independência Preferencial

X_1 and X_2 preferentially independent of X_3 iff
preference between $\langle x_1, x_2, x_3 \rangle$ and $\langle x'_1, x'_2, x_3 \rangle$
does not depend on x_3

E.g., $\langle Noise, Cost, Safety \rangle$:
⟨20,000 suffer, \$4.6 billion, 0.06 deaths/mpm⟩ vs.
⟨70,000 suffer, \$4.2 billion, 0.06 deaths/mpm⟩

$\langle Noise, Cost, Safety \rangle$:
⟨20,000 suffer, \$4.6 billion, 0.6 deaths/mpm⟩ vs.
⟨70,000 suffer, \$4.2 billion, 0.6 deaths/mpm⟩

Theorem (Leontief, 1947): if every pair of attributes is P.I. of its complement, then every subset of attributes is P.I of its complement: mutual P.I..

Theorem (Debreu, 1960): mutual P.I. $\Rightarrow$ $\exists$ additive value function:

$$V(S) = \Sigma_i V_i(X_i(S))$$

Hence assess n single-attribute functions; often a good approximation

Independência da Utilidade

Need to consider preferences over lotteries:
$\mathbf{X}$ is utility-independent of $\mathbf{Y}$ iff
preferences over lotteries $\mathbf{X}$ do not depend on $\mathbf{y}$

Mutual U.I.: each subset is U.I of its complement
$\Rightarrow$ $\exists$ multiplicative utility function:

$$U = k_1U_1 + k_2U_2 + k_3U_3$$
$$+ k_1k_2U_1U_2 + k_2k_3U_2U_3 + k_3k_1U_3U_1$$
$$+ k_1k_2k_3U_1U_2U_3$$

Routine procedures and software packages for generating preference tests to identify various canonical families of utility functions

Problemas na Teoria da maximização da utilidade esperada

- A teoria da maximização da utilidade esperada é uma teoria normativa. Ela descreve como um agente deve reagir. Entretanto, não é uma teoria descritiva da tomada de decisões reais

- Há evidências experimentais que as pessoas violam os axiomas da teoria da utilidade

Escolha A ou B

- A: 80% de chance de ganhar $4000
- B: 100% de chance de ganhar $3.000

Escolha C ou D

- C: 20% de chance de ganhar $4000
- D: 25% de chance de ganhar $3.000

Supondo U(0)=0

- Se maioria escolhe B em detrimento de A e C em detrimento de D,
  - De A e B, temos que $0,8*U(4000)<U(3000)$
  - De C e D temos que $0,8U(4000)>U(3000)$
- Contraditório!!!!

Teorias alternativas

- Em linhas gerais as pessoas divergem da teoria da maximização da utilidade esperada em situações de probabilidade muito alta e/ou muito baixa

- Há algumas teorias alternativas que se propõem a descrever o comportamento humano real. Uma das mais relevantes foi proposta por Kahneman e Tversky. Esta teoria propõe um modelo alternativo que descreve esse efeito "certeza" e outros

Utilities and Preferences for Agents

- Assume we have just two agents: $Ag = \{i, j\}$
- Agents are assumed to be *self-interested*: they *have preferences over how the environment is*
- Assume $\Omega = \{\omega_1, \omega_2, \ldots\}$ is the set of "outcomes" that agents have preferences over
- We capture preferences by *utility functions*:

$$u_i = \Omega \rightarrow \mathbb{R}$$
$$u_j = \Omega \rightarrow \mathbb{R}$$

- Utility functions lead to *preference orderings* over outcomes:

$$\omega \succeq_i \omega' \text{ means } u_i(\omega) \geq u_i(\omega')$$

-

$$\omega \succ_j \omega' \text{ means } u_j(\omega) > u_j(\omega')$$

Multiagent Encounters

- We need a model of the environment in which these agents will act…
  - agents simultaneously choose an action to perform, and as a result of the actions they select, an outcome in Ω will result
  - the *actual* outcome depends on the *combination* of actions

- Environment behavior may be given by *state transformer function*:

$$\tau : \underbrace{Ac}_{\text{agent } i\text{'s action}} \times \underbrace{Ac}_{\text{agent } j\text{'s action}} \rightarrow \Omega$$

Non-cooperative Game Theory

- What is it?
  - mathematical study of interaction between rational, self-interested agents
- Why is it called non-cooperative?
  - while it's most interested in situations where agents' interests conflict, it's not restricted to these settings
  - the key is that the individual is the basic modeling unit, and that individuals pursue their own interests
    - cooperative/coalitional game theory has teams as the central unit, rather than agents

Defining Games

- Finite, n-person game: $\langle N, A, u \rangle$:
  - N is a finite set of n players, indexed by i
  - $A = A_1 \times \ldots \times A_n$, where A_i is the action set for player i
    - $a \in A$ is an action profile, and so A is the space of action profiles
  - $u = \langle u_1, \ldots, u_n \rangle$, a utility function for each player, where $u_i : A \mapsto \mathbb{R}$

- Writing a 2-player game as a matrix:
  - row player is player 1, column player is player 2
  - rows are actions $a \in A_1$, columns are $a' \in A_2$
  - cells are outcomes, written as a tuple of utility values for each player

Normal (Strategic) Form Games

- Normal Form (Strategic Form): Outcome depends only on agent's actions

- Non-normal form: outcome may depends on environment (randomnly)

Prisioner's dilemma

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

Prisoner's dilemma

Prisoner's dilemma is any game

	C	D
C	a, a	b, c
D	c, b	d, d

with $c > a > d > b$.

Games of Pure Competition

Players have exactly opposed interests

- There must be precisely two players (otherwise they can't have exactly opposed interests)
- For all action profiles $a \in A$, $u_1(a) + u_2(a) = c$ for some constant c
  - Special case: zero sum

	Heads	Tails
Heads	1, −1	−1, 1
Tails	−1, 1	1, −1

Games of Cooperation

Players have exactly the same interests.

- no conflict: all players want the same things
- $\forall a \in A, \forall i, j,\ u_i(a) = u_j(a)$

	Left	Right
Left	1	0
Right	0	1

General Games

The most interesting games combine elements of cooperation *and* competition.

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

Analyzing games

- We've defined some canonical games, and thought about how to play them. Now let's examine the games from the outside
- From the point of view of an outside observer, can some outcomes of a game be said to be better than others?
    - we have no way of saying that one agent's interests are more important than another's
    - intuition: imagine trying to find the revenue-maximizing outcome when you don't know what currency has been used to express each agent's payoff
- Are there situations where we can still prefer one outcome to another?

Pareto Optimatility

- Idea: sometimes, one outcome o is at least as good for every agent as another outcome o', and there is some agent who strictly prefers o to o'
  - in this case, it seems reasonable to say that o is better than o'
  - we say that o Pareto-dominates o'.

- An outcome o^* is Pareto-optimal if there is no other outcome that Pareto-dominates it.
  - can a game have more than one Pareto-optimal outcome?
  - does every game have at least one Pareto-optimal outcome?

Pareto Optimatility in Examples

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

	Left	Right
Left	1	0
Right	0	1

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

	Heads	Tails
Heads	1, −1	−1, 1
Tails	−1, 1	1, −1

Pareto Optimatility and Prisioner's Dilemma

The Prisoner's Dilemma

- (C,C) is Pareto optimal
  - No profile gives both players a higher payoff
- (D,C) is Pareto optimal
  - No profile gives player 1 a higher payoff
- (D,C) is Pareto optimal - same argument
- (D,D) is Pareto dominated by (C,C)
  - But ironically, (D,D) is the dominant strategy equilibrium

Pareto Optimatility in Examples

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

	Left	Right
Left	1	0
Right	0	1

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

	Heads	Tails
Heads	1, −1	−1, 1
Tails	−1, 1	1, −1

?

Best Response and Nash Equilibrium

- If you knew what everyone else was going to do, it would be easy to pick your own action
- Let $a_{-i} = \langle a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n \rangle$.
  - now $a = (a_{-i}, a_i)$

- Best response: $a_i^* \in BR(a_{-i})$ iff
  $\forall a_i \in A_i,\ u_i(a_i^*, a_{-i}) \geq u_i(a_i, a_{-i})$

Nash Equilibrium

- Now let's return to the setting where no agent knows anything about what the others will do
- What can we say about which actions will occur?

- Idea: look for stable action profiles.
- $a = \langle a_1, \ldots, a_n \rangle$ is a ("pure strategy") Nash equilibrium iff $\forall i,\ a_i \in BR(a_{-i})$.

Nash Equilibrium in Examples

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

	Left	Right
Left	1	0
Right	0	1

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

	Heads	Tails
Heads	$1, -1$	$-1, 1$
Tails	$-1, 1$	$1, -1$

Nash Equilibria in Examples

	C	D
C	−1, −1	−4, 0
D	0, −4	**−3, −3**

	Left	Right
Left	1	0
Right	0	1

?

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

?

	Heads	Tails
Heads	1, −1	−1, 1
Tails	−1, 1	1, −1

?

The paradox of Prisoner's dilemma: the Nash equilibrium is the only non-Pareto-optimal outcome!

Mixed Strategies

- It would be a pretty bad idea to play any deterministic strategy in matching pennies
- Idea: confuse the opponent by playing randomly
- Define a strategy s_i for agent i as any probability distribution over the actions A_i.
  - pure strategy: only one action is played with positive probability
  - mixed strategy: more than one action is played with positive probability
    - these actions are called the support of the mixed strategy
- Let the set of all strategies for i be S_i
- Let the set of all strategy profiles be $S = S_1 \times \ldots \times S_n$.

Utility under Mixed Strategies

- What is your payoff if all the players follow mixed strategy profile $s \in S$?
  - We can't just read this number from the game matrix anymore: we won't always end up in the same cell
- Instead, use the idea of expected utility from decision theory:

$$u_i(s) = \sum_{a \in A} u_i(a) Pr(a|s)$$

$$Pr(a|s) = \prod_{j \in N} s_j(a_j)$$

Best Response and Nash Equilibrium

Our definitions of best response and Nash equilibrium generalize from actions to strategies.

- Best response:
  - $s_i^* \in BR(s_{-i})$ iff $\forall s_i \in S_i$, $u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$
- Nash equilibrium:
  - $s = \langle s_1, \ldots, s_n \rangle$ is a Nash equilibrium iff $\forall i$, $s_i \in BR(s_{-i})$
- Every finite game has a Nash equilibrium! [Nash, 1950]
  - e.g., matching pennies: both players play heads/tails 50%/50%

Computing Mixed Strategy: Battle of Sexes

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- It's hard in general to compute Nash equilibria, but it's easy when you can guess the support
- For BoS, let's look for an equilibrium where all actions are part of the support

Computing Mixed Strategy: Battle of Sexes

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- Let player 2 play B with p, F with $1 - p$.
- If player 1 best-responds with a mixed strategy, player 2 must make him indifferent between F and B (why?)

$$u_1(B) = u_1(F)$$
$$2p + 0(1-p) = 0p + 1(1-p)$$
$$p = \frac{1}{3}$$

Computing Mixed Strategy: Battle of Sexes

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- Likewise, player 1 must randomize to make player 2 indifferent.
  - Why is player 1 willing to randomize?
- Let player 1 play B with q, F with $1 - q$.

$$u_2(B) = u_2(F)$$
$$q + 0(1-q) = 0q + 2(1-q)$$
$$q = \frac{2}{3}$$

- Thus the mixed strategies $(\frac{2}{3}, \frac{1}{3})$, $(\frac{1}{3}, \frac{2}{3})$ are a Nash equilibrium.

Matching Pennies's Nash Equilibrium

Agent 1 \ Agent 2	Heads	Tails
Heads	1, –1	–1, 1
Tails	–1, 1	1, –1

- Each agent has a penny
- Each agent independently chooses to display his/her penny heads up or tails up
- Easy to see that in this game, no pure strategy could be part of a Nash equilibrium
  - For each combination of pure strategies, one of the agents can do better by changing his/her strategy
    - for (Heads,Heads), agent 2 can do better by switching to Tails
    - for (Heads,Tails), agent 1 can do better by switching to Tails
    - for (Tails,Tails), agent 2 can do better by switching to Heads
    - for (Tails,Heads), agent 1 can do better by switching to Heads
- But there's a mixed-strategy equilibrium:
  - (s,s), where $s(\text{Heads}) = s(\text{Tails}) = ½$

Interpreting Mixed Strategies

What does it mean to play a mixed strategy? Different interpretations:

- Randomize to confuse your opponent
  - consider the matching pennies example
- Players randomize when they are uncertain about the other's action
  - consider battle of the sexes
- Mixed strategies are a concise description of what might happen in repeated play: count of pure strategies in the limit
- Mixed strategies describe population dynamics: 2 agents chosen from a population, all having deterministic strategies. MS is the probability of getting an agent who will play one PS or another.