Reasoning in Games: Players as Programs

Lecture 1

Eric Pacuit
Department of Philosophy
University of Maryland
pacuit.org

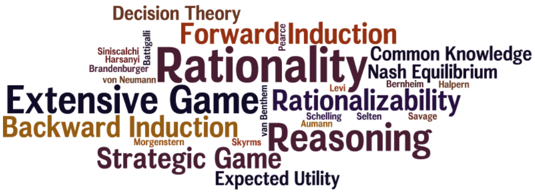

Plan

Monday Epistemic utility theory, Decision- and game-theoretic background: Nash equilibrium

Tuesday Introduction to game theory: rationalizability, epistemic game theory, forward and backward induction; Iterated games and learning, Skyrms's model of rational deliberation I

Wednesday Skyrms's model of rational deliberation II; Introduction to webppl; Game-theoretic reasoning in webppl

Thursday Coordination games (comparing Skyrms's model of deliberation and the webppl approach)

Friday Models of game-theoretic reasoning

The Guessing Game

Guess a number between 1 & 100.
The closest to 2/3 of the average wins.

pacuit.org/games/avg

The Guessing Game, again

Guess a number between 1 & 100.
The closest to 2/3 of the average wins.

The Guessing Game, again

Guess a number between 1 & 100.
The closest to 2/3 of the average wins.

pacuit.org/games/avg

The Guessing Game

Guess a number between 1 & 100.
The closest to 2/3 of the average wins.

The Guessing Game

Guess a number between 1 & 100.
The closest to 2/3 of the average wins.

What number should you guess?

The Guessing Game

Guess a number between 1 & 100.
The closest to 2/3 of the average wins.

What number should you guess? 100

The Guessing Game

Guess a number between 1 & 100.
The closest to 2/3 of the average wins.

What number should you guess? ~~100~~, 99

The Guessing Game

Guess a number between 1 & 100.
The closest to 2/3 of the average wins.

What number should you guess? ~~100~~, ~~99~~, . . . , 67

The Guessing Game

Guess a number between 1 & 100.
The closest to 2/3 of the average wins.

What number should you guess? ~~100~~, ~~99~~, …, ~~67~~,…, 2, 1

The Guessing Game

Guess a number between 1 & 100.
The closest to 2/3 of the average wins.

What number should you guess? ~~100~~, ~~99~~, . . . , ~~67~~, . . . , ~~2~~, 1

Traveler's Dilemma

1. You and your friend write down an integer between 2 and 100 (without discussing).

Traveler's Dilemma

1. You and your friend write down an integer between 2 and 100 (without discussing).
2. If both of you write down the same number, then both will receive that amount in dollars from the airline in compensation.

Traveler’s Dilemma

1. You and your friend write down an integer between 2 and 100 (without discussing).
2. If both of you write down the same number, then both will receive that amount in dollars from the airline in compensation.
3. If the numbers are different, then the airline assumes that the smaller number is the actual price of the luggage.

Traveler’s Dilemma

1. You and your friend write down an integer between 2 and 100 (without discussing).
2. If both of you write down the same number, then both will receive that amount in dollars from the airline in compensation.
3. If the numbers are different, then the airline assumes that the smaller number is the actual price of the luggage.
4. The person that wrote the smaller number will receive that amount plus $2 (as a reward), and the person that wrote the larger number will receive the smaller number minus $2 (as a punishment).

Traveler’s Dilemma

1. You and your friend write down an integer between 2 and 100 (without discussing).
2. If both of you write down the same number, then both will receive that amount in dollars from the airline in compensation.
3. If the numbers are different, then the airline assumes that the smaller number is the actual price of the luggage.
4. The person that wrote the smaller number will receive that amount plus $2 (as a reward), and the person that wrote the larger number will receive the smaller number minus $2 (as a punishment).

Suppose that you are randomly paired with another person from class. What number would you write down?

pacuit.org/games/td

Plan

- Expected utility
- Basic game-theoretic reasoning: Nash equilibrium, rationalizability
- Epistemic game theory, correlated equilibrium
- Backward and forward induction

Decision Problems

In many circumstances the decision maker doesn't get to choose outcomes directly, but rather chooses an instrument that affects what outcome actually occurs.

Decision Problems

In many circumstances the decision maker doesn't get to choose outcomes directly, but rather chooses an instrument that affects what outcome actually occurs.

Economists distinguish between choice under:

- *certainty*: highly confident about the relationship between actions and outcomes
- *risk*: clear sense of possibilities and their likelihoods
- *uncertainty*: the relationship between actions and outcomes is so imprecise that it is not possible to assign likelihoods

Decision Problems

A

B

Decision Problems

	w_1	w_2	$\cdots$	w_{n-1}	w_n
A					
B					

Decision Problems

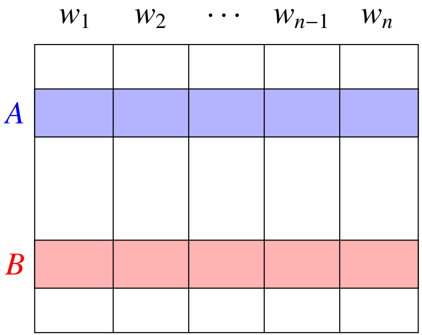

An **act** is a function $F : W \to O$

Strict Dominance

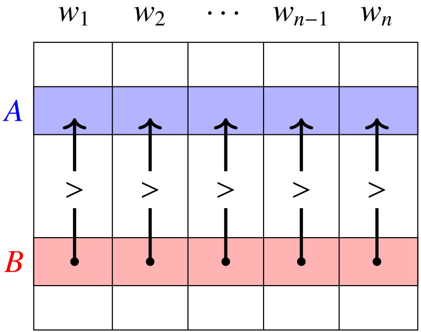

$$\forall\, w \in W,\ u(A(w)) > u(B(w))$$

Weak Dominance

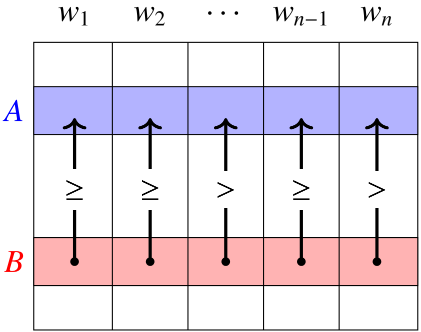

$$\forall\, w \in W,\ u(A(w)) \geq u(B(w)) \text{ and } \exists\, w \in W,\ u(A(w)) > u(B(w))$$

MaxMin (Security)

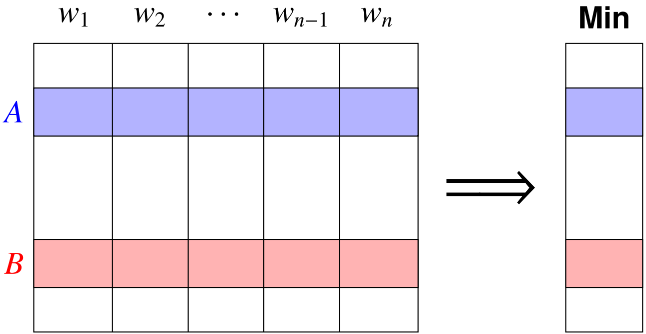

$$\min(\{u(A(w)) \mid w \in W\}) > \min(\{u(B(w)) \mid w \in W\})$$

MaxMax

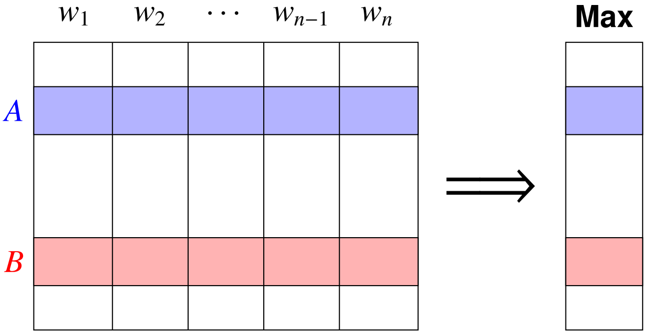

$$\max(\{u(A(w)) \mid w \in W\}) > \max(\{u(B(w)) \mid w \in W\})$$

Maximize (Subjective) Expected Utility

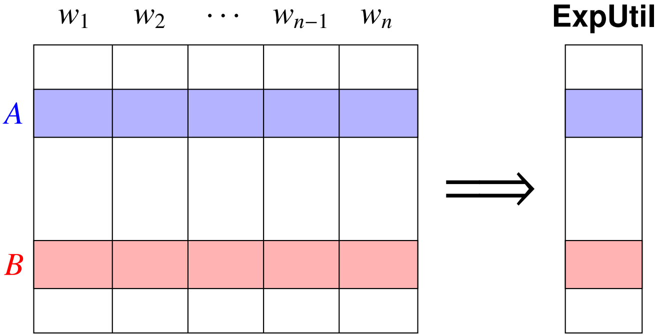

$$\sum_{w \in W} P_A(w) * u(A(w)) > \sum_{w \in W} P_A(w) * u(B(w))$$

Subjective Expected Utility

Probability: Suppose that $W = \{w_1, \ldots, w_n\}$ is a finite set of states. A probability function on W is a function $P : W \rightarrow [0, 1]$ where $\sum_{w \in W} P(w) = 1$ (i.e., $P(w_1) + P(w_2) + \cdots + P(w_n) = 1$).

Suppose that A is an act for a set of outcomes O (i.e., $A : W \rightarrow O$). The **expected utility** of A is:

$$\sum_{w \in W} P(w) * u(A(w))$$

$$EU(A) = \sum_{o \in O} P_A(o) \times U(o)$$

$$EU(A) = \sum_{o \in O} P_A(o) \times U(o)$$

Expected utility of action A

$$EU(A) = \sum_{o \in O} P_A(o) \times U(o)$$

Expected utility of action A

Utility of outcome o

$$EU(A) = \sum_{o \in O} P_A(o) \times U(o)$$

Expected utility of action A

Utility of outcome o

Probability of outcome o conditional on A

$P_A(o)$: probability of o conditional on A — how likely it is that outcome o will occur, on the supposition that the agent chooses act A.

$P_A(o)$: probability of o conditional on A — how likely it is that outcome o will occur, on the supposition that the agent chooses act A.

Evidential: $P_A(o) = P(o \mid A) = \frac{P(o \ \& \ A)}{P(A)}$

$P_A(o)$: probability of o conditional on A — how likely it is that outcome o will occur, on the supposition that the agent chooses act A.

Evidential: $P_A(o) = P(o \mid A) = \frac{P(o \ \& \ A)}{P(A)}$

Classical: $P_A(o) = \sum_{s \in S} P(s) f_{A,s}(o)$, where

$$f_{A,s}(o) = \begin{cases} 1 & A(s) = o \\ 0 & A(s) \neq o \end{cases}$$

$P_A(o)$: probability of o conditional on A — how likely it is that outcome o will occur, on the supposition that the agent chooses act A.

Evidential: $P_A(o) = P(o \mid A) = \frac{P(o \ \& \ A)}{P(A)}$

Classical: $P_A(o) = \sum_{s \in S} P(s) f_{A,s}(o)$, where

$$f_{A,s}(o) = \begin{cases} 1 & A(s) = o \\ 0 & A(s) \neq o \end{cases}$$

Causal: $P_A(o) = P(A \mathrel{\Box\!\!\rightarrow} o)$

P(“if A were performed, outcome o would ensue”)

(Lewis, 1981)

Dominance and Act-State Dependence

	w_1	w_2
A	1	3
B	2	4

Dominance and Act-State Dependence

	w_1	w_2
A	1	3
B	2	4

Dominance and Act-State Dependence

Dominance reasoning is appropriate only when probability of outcome is *independent of choice*.

(A nasty nephew wants inheritance from his rich Aunt. The nephew wants the inheritance, but other things being equal, does not want to apologize. Does dominance give the nephew a reason to not apologize? *Whether or not the nephew is cut from the will may depend on whether or not he apologizes*.)

Why maximize expected utility?

Why maximize expected utility?

Law of Large Numbers: everyone who maximizes expected utility will *almost certainly* be better off in the long run. By performing a random experiment sufficiently many times, the probability that the average outcome differs from the expected outcome can be rendered *arbitrarily* small.

Why maximize expected utility?

Law of Large Numbers: everyone who maximizes expected utility will *almost certainly* be better off in the long run. By performing a random experiment sufficiently many times, the probability that the average outcome differs from the expected outcome can be rendered *arbitrarily* small.

Gambler's Ruin: Suppose Ann and Bob start with $1000 each and flip a fair coin.

Why maximize expected utility?

Law of Large Numbers: everyone who maximizes expected utility will *almost certainly* be better off in the long run. By performing a random experiment sufficiently many times, the probability that the average outcome differs from the expected outcome can be rendered *arbitrarily* small.

Gambler's Ruin: Suppose Ann and Bob start with \$1000 each and flip a fair coin. Ann gives Bob \$1 if H and Bob gives Ann \$1 if T.

Why maximize expected utility?

Law of Large Numbers: everyone who maximizes expected utility will *almost certainly* be better off in the long run. By performing a random experiment sufficiently many times, the probability that the average outcome differs from the expected outcome can be rendered *arbitrarily* small.

Gambler's Ruin: Suppose Ann and Bob start with \$1000 each and flip a fair coin. Ann gives Bob \$1 if H and Bob gives Ann \$1 if T. If they flip the coin a *sufficiently* large number of times, each player is *guaranteed* to face a sequence of flips that bankrupts them.

Why maximize expected utility?

Law of Large Numbers: everyone who maximizes expected utility will *almost certainly* be better off in the long run. By performing a random experiment sufficiently many times, the probability that the average outcome differs from the expected outcome can be rendered *arbitrarily* small.

Gambler's Ruin: Suppose Ann and Bob start with \$1000 each and flip a fair coin. Ann gives Bob \$1 if H and Bob gives Ann \$1 if T. If they flip the coin a *sufficiently* large number of times, each player is *guaranteed* to face a sequence of flips that bankrupts them. The player that faces such a sequence first, will never have an opportunity to feel the effects of the Law of Large Numbers.

Expected Utility Theory

R. A. Briggs. *Normative Theories of Rational Choice: Expected Utility*. Stanford Encyclopedia of Philosophy, 2014
`https://plato.stanford.edu/entries/rationality-normative-utility/`.

Expected Utility Theory

Representation theorems by Von Neumann-Morgenstern, Aumann-Anscombe, and Savage.

Expected Utility Theory

Representation theorems by Von Neumann-Morgenstern, Aumann-Anscombe, and Savage. Some issues:

- The axioms are too strong. Do rational decisions *have* to obey these axioms?

Expected Utility Theory

Representation theorems by Von Neumann-Morgenstern, Aumann-Anscombe, and Savage. Some issues:

- The axioms are too strong. Do rational decisions *have* to obey these axioms?
- No action guidance. Rational decision makers do not prefer an act *because* its expected utility is favorable, but can only be described as *if* they were acting from this principle.

Expected Utility Theory

Representation theorems by Von Neumann-Morgenstern, Aumann-Anscombe, and Savage. Some issues:

- The axioms are too strong. Do rational decisions *have* to obey these axioms?
- No action guidance. Rational decision makers do not prefer an act *because* its expected utility is favorable, but can only be described as *if* they were acting from this principle.
- Utility without chance. It seems odd from a linguistic point of view to say that the *meaning* of utility has something to do with preferences over lotteries.

Issues

- Preference, choice and utility
- Preferences satisfy completeness and transitivity (Money-pump argument)
- Allais paradox: risk-aversion
- Ellsberg paradox: ambiguity-aversion
- Newcomb's paradox, Death in Damascus, Pyschopath button problem, irrational choice: Act-state dependence
- Framing issues

From Decisions to Games, I

Commenting on the difference between Robin Crusoe's maximization problem and the maximization problem faced by participants in a social economy, von Neumann and Morgenstern write:

"Every participant can determine the variables which describe his own actions but not those of the others. Nevertheless those "alien" variables cannot, from his point of view, be described by statistical assumptions.

From Decisions to Games, I

Commenting on the difference between Robin Crusoe's maximization problem and the maximization problem faced by participants in a social economy, von Neumann and Morgenstern write:

"Every participant can determine the variables which describe his own actions but not those of the others. Nevertheless those "alien" variables cannot, from his point of view, be described by statistical assumptions. This is because the others are guided, just as he himself, by rational principles—whatever that may mean—and no *modus procedendi* can be correct which does not attempt to understand those principles and the interactions of the conflicting interests of all participants."

(vNM, pg. 11)

Game Situations

1. a group of *self-interested* agents (players) involved in some interdependent decision problem

Game Situations

	Bob	
L		R
1		0

1. a group of *self-interested* agents (players) involved in some interdependent decision problem

Game Situations

	Bob	
	L	R
	1	0
	0	1

1. a group of *self-interested* agents (players) involved in some interdependent decision problem

Game Situations

		Bob	
		L	R
Ann	U	1 1	0 0
	D	0 0	1 1

1. a group of *self-interested* agents (players) involved in some interdependent decision problem

Game Situations

		Bob	
		L	R
Ann	U	1,1	0,0
	D	0,0	1,1

1. a group of *self-interested* agents (players) involved in some interdependent decision problem

Game Situations

		Bob	
		L	R
Ann	U	1,1	0,0
	D	0,0	1,1

1. a group of *self-interested* agents (players) involved in some interdependent decision problem

Just Enough Game Theory

A **game** is a mathematical model of a strategic interaction that includes

- the actions the players *can* take
- the players' interests (i.e., preferences),
- the "structure" of the decision problem

Just Enough Game Theory

A **game** is a mathematical model of a strategic interaction that includes

- the actions the players *can* take
- the players' interests (i.e., preferences),
- the "structure" of the decision problem

It does not specify the actions that the players do take.

Games

		B	
		l	r
A	u	3, 1	0, 0
	d	0, 0	1, 3

Games

		B	
		l	r
A	u	3, 1	0, 0
	d	0, 0	1, 3

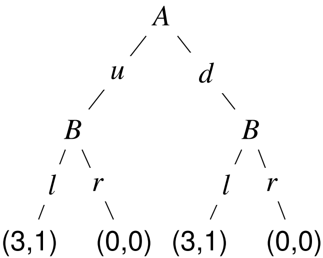

Games

		B	
		l	r
A	u	3, 1	0, 0
	d	0, 0	1, 3

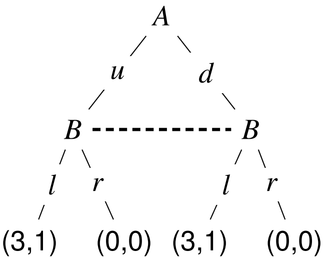

Games

		B	
		l	r
A	u	3, 1	0, 0
	d	0, 0	1, 3

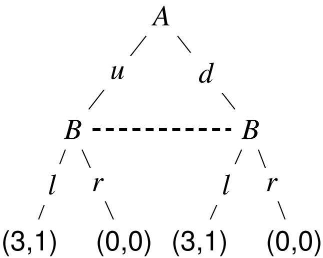

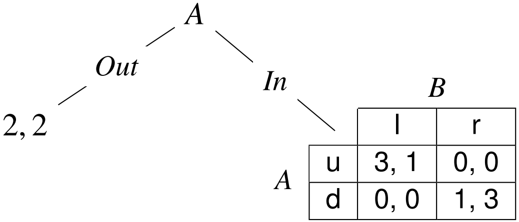

From Decisions to Games, II

"*[T]he* fundamental insight of game theory [is] that a rational player must take into account that the players reason about each other in deciding how to play."

R. Aumann and J. Dreze. *Rational Expectations in Games*. American Economic Review, 98, pp. 72-86, 2008.

Solution Concept

A **solution concept** is a systematic description of the outcomes that may emerge in a family of games.

This is the starting point for most of game theory and includes many variants: Nash equilibrium, backwards induction, or iterated dominance of various kinds.

These are usually thought of as the embodiment of “rational behavior” in some way and used to analyze game situations.

Let $G = \langle (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$ be a finite strategic game (each S_i is finite and the set of players N is finite).

A **strategy profile** is an element $\sigma \in S = S_1 \times \cdots \times S_n$

σ is a (pure strategy) **Nash equilibrium** provided for all i, for all $s_i \in S_i$,

$$u_i(\sigma) \geq u_i(s_i, \sigma_{-i})$$

Zero-Sum Games

		Bob	
		L	R
Ann	U	1,4	4,1
	D	2,3	3,2

What should Ann *do*?

Zero-Sum Games

		Bob	
		L	R
Ann	U	1,4	4,1
	D	2,3	3,2

What should Ann *do*? *Bob best choice in Ann's worst choice*

Zero-Sum Games

		Bob L	Bob R	
Ann	U	1,4	4,1	1
Ann	D	2,3	3,2	2

What should Ann *do*? *Security strategy: minimize over each row and choose the maximum value*

Zero-Sum Games

		Bob	
		L	R
Ann	U	1,4	4,1
	D	2,3	3,2
		3	1

What should Bob *do*? *Security strategy: minimize over each column and choose the maximum value*

Zero-Sum Games

		Bob	
		L	R
Ann	U	1,4	4,1
	D	2,3	3,2

The profile of security strategies (D, L) is a Nash equilibrium

Matching Pennies

		Bob H	Bob T
Ann	H	1,-1	-1, 1
Ann	T	-1,1	1,-1

There are no pure strategy Nash equilibria.

Mixed Strategies

		Bob	
		H	T
Ann	H	1,-1	-1, 1
	T	-1,1	1,-1

A **mixed strategy** is a probability distribution over the set of pure strategies. For instance:

- $[1/2 : H, 1/2 : T]$
- $[1/3 : H, 2/3 : T]$
- ...

Mixed Extension

	L	R
U	$1, -1$	$-1, 1$
D	$-1, 1$	$1, -1$

Mixed Extension

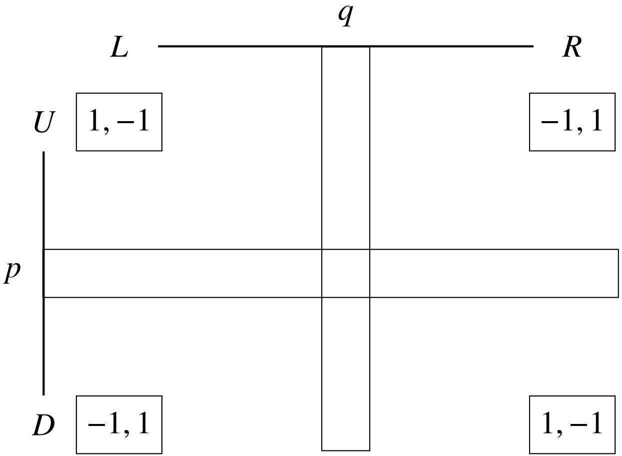

Mixed Extension

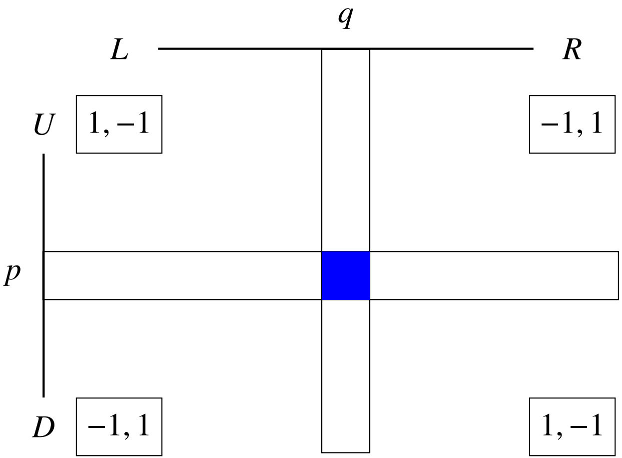

$$pq - p(1-q) - (1-p)q + (1-p)(1-q),\ -pq + p(1-q) + (1-p)q - (1-p)(1-q)$$

Matching Pennies

		Bob	
		H	T
Ann	H	1,-1	-1, 1
	T	-1,1	1,-1

The mixed strategy $([1/2 : H, 1/2 : T], [1/2 : H, 1/2 : T])$ is the only Nash equilibrium.

Theorem (Von Neumann). For every two-player zero- sum game with finite strategy sets S_1 and S_2, there is a number v, called the **value** of the game such that:

1. $v = \max_{p \in \Delta(S_1)} \min_{q \in \Delta(S_2)} U_1(p, q) = \min_{q \in \Delta(S_2)} \max_{p \in \Delta(S_1)} U_1(p, q)$
2. The set of mixed Nash equilibria is nonempty. A mixed strategy profile (p, q) is a Nash equilibrium if and only if

$$p \in \text{argmax}_{p \in \Delta(S_1)} \min_{q \in \Delta(S_2)} U_1(p, q)$$

$$q \in \text{argmax}_{q \in \Delta(S_2)} \min_{p \in \Delta(S_1)} U_1(p, q)$$

3. For all mixed Nash equilibria (p, q), $U_1(p, q) = v$

In zero-sum games

- There exists a mixed strategy Nash equilibrium
- There may be more than one Nash equilibria
- Security strategies are always a Nash equilibrium
- Components of Nash equilibria are interchangeable: If σ and σ' are Nash equilibria in a 2-player game, then (σ_1, σ_2') is also a Nash equilibrium.

Let $G = \langle (S_i)_{i\in N}, (u_i)_{i\in N} \rangle$ be a finite strategic game.

$\Sigma_i = \{p \mid p : S_i \to [0, 1]$ and $\sum_{s_i \in S_i} p(s_i) = 1\}$

The **mixed extension** of G is the game $\langle (\Sigma_i)_{i\in N}, (U_i)_{i\in N} \rangle$ where for $\sigma \in \Sigma = \Sigma_1 \times \cdots \times \Sigma_n$:

$$U_i(\sigma) = \sum_{(s_1,\ldots,s_n)\in S} \sigma_1(s_1)\sigma_2(s_2)\cdots\sigma_n(s_n)u_i(s_1, \ldots, s_n)$$

Theorem (Nash). Every finite game G has a Nash equilibrium in mixed strategies (i.e., there is a Nash equilibrium in the mixed extension G).

Mixed Strategies

"We are reluctant to believe that our decisions are made at random. We prefer to be able to point to a reason for each action we take. Outside of Las Vegas we do not spin roulettes." (Rubinstein)

Mixed Strategies

"We are reluctant to believe that our decisions are made at random. We prefer to be able to point to a reason for each action we take. Outside of Las Vegas we do not spin roulettes." (Rubinstein)

- One can think about a game as an interaction between large populations...a mixed strategy is viewed as the distribution of the pure choices in the population.

Mixed Strategies

"We are reluctant to believe that our decisions are made at random. We prefer to be able to point to a reason for each action we take. Outside of Las Vegas we do not spin roulettes." (Rubinstein)

- One can think about a game as an interaction between large populations...a mixed strategy is viewed as the distribution of the pure choices in the population.
- *Harsanyi's purification theorem*: A player's mixed strategy is thought of as a plan of action which is dependent on private information which is not specified in the model. Although the player's behavior appears to be random, it is actually deterministic.

Mixed Strategies

"We are reluctant to believe that our decisions are made at random. We prefer to be able to point to a reason for each action we take. Outside of Las Vegas we do not spin roulettes." (Rubinstein)

- One can think about a game as an interaction between large populations...a mixed strategy is viewed as the distribution of the pure choices in the population.
- *Harsanyi's purification theorem*: A player's mixed strategy is thought of as a plan of action which is dependent on private information which is not specified in the model. Although the player's behavior appears to be random, it is actually deterministic.
- Mixed strategies are beliefs held by all *other* players concerning a player's actions.