

Normal-form games


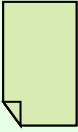


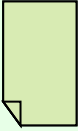

Vincent Conitzer


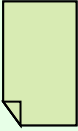

“2/3 of the average” game

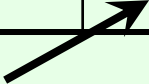
- Everyone writes down a number between 0 and 100
- Person closest to $2/3$ of the average wins
- Example:
 - A says 50
 - B says 10
 - C says 90
 - Average(50, 10, 90) = 50
 - $2/3$ of average = 33.33
 - A is closest ($|50-33.33| = 16.67$), so A wins

Rock-paper-scissors

Column player aka.
player 2
(simultaneously)
chooses a column



	0, 0	-1, 1	1, -1
	1, -1	0, 0	-1, 1
	-1, 1	1, -1	0, 0




Row player
aka. player 1
chooses a row

A row or column is
called an **action** or
(pure) strategy

Row player's utility is always listed first, column player's second


Zero-sum game: the utilities in each entry sum to 0 (or a constant)
Three-player game would be a 3D table with 3 utilities per entry, etc.

Matching pennies (~penalty kick)



L

R



L

1, -1

-1, 1


R

-1, 1

1, -1

“Chicken”


- Two players drive cars towards each other
- If one player goes straight, that player wins
- If both go straight, they both die




	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5


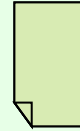




not zero-sum

Rock-paper-scissors – Seinfeld variant



MICKEY: All right, rock beats paper!
(Mickey smacks Kramer's hand for losing)
KRAMER: I thought paper covered rock.
MICKEY: Nah, rock flies right through paper.
KRAMER: What beats rock?
MICKEY: (looks at hand) Nothing beats rock.


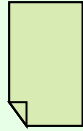


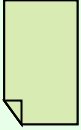



			
	0, 0	1, -1	1, -1
	-1, 1	0, 0	-1, 1
	-1, 1	1, -1	0, 0

Dominance

- Player i 's strategy s_i **strictly dominates** s_i' if
 - for any s_{-i} , $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$
- s_i **weakly dominates** s_i' if
 - for any s_{-i} , $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$; and
 - for some s_{-i} , $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$

$-i$ = "the player(s) other than i "

			
	0, 0	1, -1	1, -1
	-1, 1	0, 0	-1, 1
	-1, 1	1, -1	0, 0

Prisoner's Dilemma

- Pair of criminals has been caught
- District attorney has evidence to convict them of a minor crime (1 year in jail); knows that they committed a major crime together (3 years in jail) but cannot prove it
- Offers them a deal:
 - If both confess to the major crime, they each get a 1 year reduction
 - If only one confesses, that one gets 3 years reduction

	confess	don't confess
confess	-2, -2	0, -3
don't confess	-3, 0	-1, -1

“Should I buy an SUV?”

purchasing cost



cost: 5



cost: 3

accident cost

cost: 5



cost: 5

cost: 8







cost: 2

cost: 5


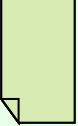



cost: 5



		
	-10, -10	-7, -11
	-11, -7	-8, -8

Mixed strategies

- **Mixed strategy** for player i = **probability distribution** over player i 's (pure) strategies
- E.g., $1/3$  , $1/3$  , $1/3$ 
- Example of dominance by a mixed strategy:

$1/2$	3, 0	0, 0
$1/2$	0, 0	3, 0
	1, 0	1, 0

A blue bracket on the left side of the table groups the first two rows, with a curved arrow pointing from the bracket to the third row, indicating that the mixed strategy of the first two rows dominates the third row.


Usage:
 σ_i denotes a mixed strategy,
 s_i denotes a pure strategy

Checking for dominance by mixed strategies


- Linear program for checking whether strategy s_i^* is **strictly** dominated by a mixed strategy:
 - maximize ε
 - such that:
 - for any s_{-i} , $\sum_{s_i} \mathbf{p}_{s_i} u_i(s_i, s_{-i}) \geq u_i(s_i^*, s_{-i}) + \varepsilon$
 - $\sum_{s_i} \mathbf{p}_{s_i} = 1$
- Linear program for checking whether strategy s_i^* is **weakly** dominated by a mixed strategy:
 - maximize $\sum_{s_{-i}} [(\sum_{s_i} \mathbf{p}_{s_i} u_i(s_i, s_{-i})) - u_i(s_i^*, s_{-i})]$
 - such that:
 - for any s_{-i} , $\sum_{s_i} \mathbf{p}_{s_i} u_i(s_i, s_{-i}) \geq u_i(s_i^*, s_{-i})$
 - $\sum_{s_i} \mathbf{p}_{s_i} = 1$

Iterated dominance

- Iterated dominance: remove (strictly/weakly) dominated strategy, repeat
- Iterated strict dominance on Seinfeld's RPS:




“2/3 of the average” game revisited



Iterated dominance: path (in)dependence

Iterated weak dominance is **path-dependent**:
sequence of eliminations may determine which
solution we get (if any)
(whether or not dominance by mixed strategies allowed)



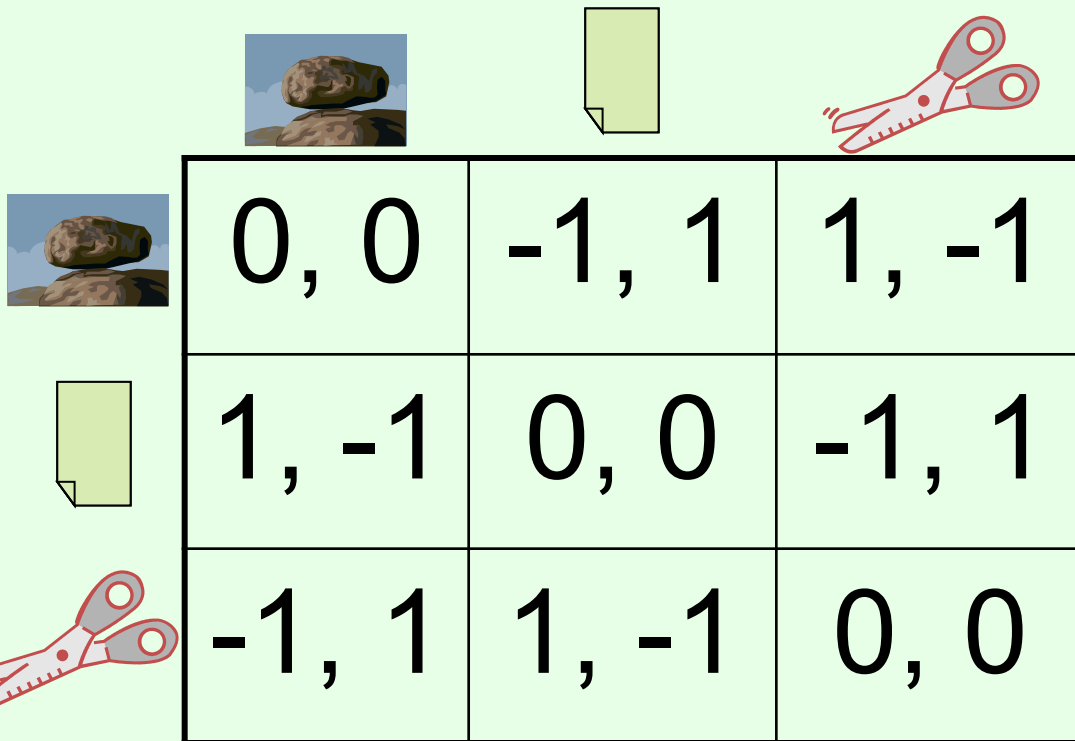
Iterated strict dominance is **path-independent**: elimination
process will always terminate at the same point
(whether or not dominance by mixed strategies allowed)



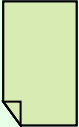

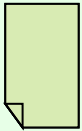

Two computational questions for iterated dominance

- 1. Can a **given strategy** be eliminated using iterated dominance?
- 2. Is there some path of elimination by iterated dominance such that only **one strategy per player remains**?
- For strict dominance (with or without dominance by mixed strategies), both can be solved in polynomial time due to path-independence:
 - Check if any strategy is dominated, remove it, repeat
- For weak dominance, both questions are NP-hard (even when all utilities are 0 or 1), with or without dominance by mixed strategies [Conitzer, Sandholm 05]
 - Weaker version proved by [Gilboa, Kalai, Zemel 93]

Two-player zero-sum games revisited

- Recall: in a zero-sum game, payoffs in each entry sum to zero
 - ... or to a constant: recall that we can subtract a constant from anyone's utility function without affecting their behavior
- What the one player gains, the other player loses



			
0, 0	-1, 1	1, -1	
	1, -1	0, 0	-1, 1
	-1, 1	1, -1	0, 0

Note: a general-sum k -player game can be modeled as a zero-sum $(k+1)$ -player game by adding a dummy player absorbing the remaining utility, so zero-sum games with 3 or more players have to deal with the difficulties of general-sum games; this is why we focus on 2-player zero-sum games here.

Best-response strategies

- Suppose you know your opponent's mixed strategy
 - E.g., your opponent plays rock 50% of the time and scissors 50%
- What is the best strategy for you to play?
- Rock gives $.5*0 + .5*1 = .5$
- Paper gives $.5*1 + .5*(-1) = 0$
- Scissors gives $.5*(-1) + .5*0 = -.5$
- So the best response to this opponent strategy is to (always) play rock
- There is always some **pure** strategy that is a best response
 - Suppose you have a mixed strategy that is a best response; then every one of the pure strategies that that mixed strategy places positive probability on must also be a best response

How to play matching pennies

		<i>Them</i>	
		L	R
<i>Us</i>	L	1, -1	-1, 1
	R	-1, 1	1, -1

- Assume opponent **knows our mixed strategy**
- If we play L 60%, R 40%...
- ... opponent will play R...
- ... we get $.6*(-1) + .4*(1) = -.2$
- What's optimal for us? What about rock-paper-scissors?

Matching pennies with a sensitive target

		<i>Them</i>	
		L	R
<i>Us</i>	L	1, -1	-1, 1
	R	-2, 2	1, -1

- If we play 50% L, 50% R, opponent will attack L
 - We get $.5*(1) + .5*(-2) = -.5$
- What if we play 55% L, 45% R?
- Opponent has choice between
 - L: gives them $.55*(-1) + .45*(2) = .35$
 - R: gives them $.55*(1) + .45*(-1) = .1$
- We get $-.35 > -.5$

Matching pennies with a sensitive target

		<i>Them</i>	
		L	R
<i>Us</i>	L	1, -1	-1, 1
	R	-2, 2	1, -1

- What if we play 60% L, 40% R?
- Opponent has choice between
 - L: gives them $.6*(-1) + .4*(2) = .2$
 - R: gives them $.6*(1) + .4*(-1) = .2$
- We get -.2 either way
- This is the **maximin** strategy
 - Maximizes our minimum utility

Let's change roles

		<i>Them</i>	
		L	R
<i>Us</i>	L	1, -1	-1, 1
	R	-2, 2	1, -1

- Suppose **we** know **their** strategy
- If they play 50% L, 50% R,
 - We play L, we get $.5*(1)+.5*(-1) = 0$
- If they play 40% L, 60% R,
 - If we play L, we get $.4*(1)+.6*(-1) = -.2$
 - If we play R, we get $.4*(-2)+.6*(1) = -.2$
- This is the **minimax** strategy

von Neumann's minimax theorem [1928]: maximin value = minimax value (~LP duality)

Minimax theorem [von Neumann 1928]

- Maximin utility: $\max_{\sigma_i} \min_{s_{-i}} u_i(\sigma_i, s_{-i})$
(= - $\min_{\sigma_i} \max_{s_{-i}} u_{-i}(\sigma_i, s_{-i})$)
- Minimax utility: $\min_{\sigma_{-i}} \max_{s_i} u_i(s_i, \sigma_{-i})$
(= - $\max_{\sigma_{-i}} \min_{s_i} u_{-i}(s_i, \sigma_{-i})$)
- Minimax theorem:
$$\max_{\sigma_i} \min_{s_{-i}} u_i(\sigma_i, s_{-i}) = \min_{\sigma_{-i}} \max_{s_i} u_i(s_i, \sigma_{-i})$$
- Minimax theorem does not hold with pure strategies only (example?)

Practice games

20, -20	0, 0
0, 0	10, -10

20, -20	0, 0	10, -10
0, 0	10, -10	8, -8

Solving for minimax strategies using linear programming

- maximize u_i
- subject to
 - for any s_{-i} , $\sum_{s_i} p_{s_i} u_i(s_i, s_{-i}) \geq u_i$
 - $\sum_{s_i} p_{s_i} = 1$

Can also convert linear programs to two-player zero-sum games, so they are equivalent

General-sum games

- You could still play a minimax strategy in general-sum games
 - I.e., pretend that the opponent is only trying to hurt you
- But this is not rational:

0, 0	3, 1
1, 0	2, 1

- If Column was trying to hurt Row, Column would play Left, so Row should play Down
- In reality, Column will play Right (strictly dominant), so Row should play Up
- Is there a better generalization of minimax strategies in zero-sum games to general-sum games?



Nash equilibrium

[Nash 50]



- A vector of strategies (one for each player) is called a **strategy profile**
- A strategy profile $(\sigma_1, \sigma_2, \dots, \sigma_n)$ is a **Nash equilibrium** if each σ_i is a best response to σ_{-i}
 - That is, for any i , for any σ_i' , $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i})$
- Note that this does not say anything about multiple agents changing their strategies at the same time
- In any (finite) game, at least one Nash equilibrium (possibly using mixed strategies) exists [Nash 50]
- (Note - singular: equilibrium, plural: equilibria)

Nash equilibria of “chicken”



	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

- (D, S) and (S, D) are Nash equilibria
 - They are **pure-strategy Nash equilibria**: nobody randomizes
 - They are also **strict Nash equilibria**: changing your strategy will make you strictly worse off
- No other pure-strategy Nash equilibria

Nash equilibria of “chicken” ...

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

- Is there a Nash equilibrium that uses mixed strategies? Say, where player 1 uses a mixed strategy?
- Recall: if a mixed strategy is a best response, then all of the pure strategies that it randomizes over must also be best responses
- So we need to make player 1 **indifferent** between D and S
- Player 1's utility for playing D = $-p^c_S$
- Player 1's utility for playing S = $p^c_D - 5p^c_S = 1 - 6p^c_S$
- So we need $-p^c_S = 1 - 6p^c_S$ which means $p^c_S = 1/5$
- Then, player 2 needs to be indifferent as well
- Mixed-strategy Nash equilibrium: ((4/5 D, 1/5 S), (4/5 D, 1/5 S))
 - People may die! Expected utility -1/5 for each player

The presentation game

Presenter

*Put effort into
presentation (E)*

*Do not put effort into
presentation (NE)*

*Pay
attention (A)*

4, 4

-16, -14

*Do not pay
attention (NA)*

0, -2

0, 0

- Pure-strategy Nash equilibria: (A, E), (NA, NE)
- Mixed-strategy Nash equilibrium:
((1/10 A, 9/10 NA), (4/5 E, 1/5 NE))
 - Utility 0 for audience, -14/10 for presenter
 - Can see that some equilibria are strictly better for **both** players than other equilibria, i.e., some equilibria **Pareto-dominate** other equilibria

The “equilibrium selection problem”

- You are about to play a game that you have never played before with a person that you have never met
- According to which equilibrium should you play?
- Possible answers:
 - Equilibrium that maximizes the sum of utilities (**social welfare**)
 - Or, at least not a Pareto-dominated equilibrium
 - So-called **focal** equilibria
 - “Meet in Paris” game - you and a friend were supposed to meet in Paris at noon on Sunday, but you forgot to discuss where and you cannot communicate. All you care about is meeting your friend. Where will you go?
 - Equilibrium that is the convergence point of some learning process
 - An equilibrium that is easy to compute
 - ...
- Equilibrium selection is a difficult problem

Some properties of Nash equilibria

- If you can eliminate a strategy using strict dominance or even iterated strict dominance, it will not occur in any (i.e., it will be played with probability 0 in every) Nash equilibrium
 - Weakly dominated strategies may still be played in some Nash equilibrium
- In 2-player zero-sum games, a profile is a Nash equilibrium if and only if both players play minimax strategies
 - Hence, in such games, if (σ_1, σ_2) and (σ_1', σ_2') are Nash equilibria, then so are (σ_1, σ_2') and (σ_1', σ_2)
 - No equilibrium selection problem here!

How hard is it to compute *one* (any) Nash equilibrium?

- Complexity was open for a long time
 - [Papadimitriou STOC01]: “together with factoring [...] the most important concrete open question on the boundary of P today”
- Recent sequence of papers shows that computing one (any) Nash equilibrium is PPAD-complete (even in 2-player games) [Daskalakis, Goldberg, Papadimitriou 2006; Chen, Deng 2006]
- All known algorithms require exponential time (in the worst case)

What if we want to compute a Nash equilibrium with a specific property?

- For example:
 - An equilibrium that is not Pareto-dominated
 - An equilibrium that maximizes the expected social welfare (= the sum of the agents' utilities)
 - An equilibrium that maximizes the expected utility of a given player
 - An equilibrium that maximizes the expected utility of the worst-off player
 - An equilibrium in which a given pure strategy is played with positive probability
 - An equilibrium in which a given pure strategy is played with zero probability
 - ...
- All of these are NP-hard (and the optimization questions are inapproximable assuming $P \neq NP$), even in 2-player games
[Gilboa, Zemel 89; Conitzer & Sandholm IJCAI-03/GEB-08]

Search-based approaches (for 2 players)

- Suppose we know the **support** X_i of each player i 's mixed strategy in equilibrium
 - That is, which pure strategies receive positive probability
- Then, we have a linear feasibility problem:
 - for both i , for any $s_i \in S_i - X_i$, $p_i(s_i) = 0$
 - for both i , for any $s_i \in X_i$, $\sum p_{-i}(s_{-i})u_i(s_i, s_{-i}) = u_i$
 - for both i , for any $s_i \in S_i - X_i$, $\sum p_{-i}(s_{-i})u_i(s_i, s_{-i}) \leq u_i$
- Thus, we can search over possible supports
 - This is the basic idea underlying methods in
[Dickhaut & Kaplan 91; Porter, Nudelman, Shoham AAI04/GEB08]
- Dominated strategies can be eliminated

Solving for a Nash equilibrium using MIP (2 players)

[Sandholm, Gilpin, Conitzer AAAI05]

- maximize *whatever you like (e.g., social welfare)*
- subject to
 - for both i , for any s_i , $\sum_{s_{-i}} \mathbf{p}_{s_{-i}} u_i(s_i, s_{-i}) = \mathbf{u}_{s_i}$
 - for both i , for any s_i , $\mathbf{u}_i \geq \mathbf{u}_{s_i}$
 - for both i , for any s_i , $\mathbf{p}_{s_i} \leq \mathbf{b}_{s_i}$
 - for both i , for any s_i , $\mathbf{u}_i - \mathbf{u}_{s_i} \leq M(1 - \mathbf{b}_{s_i})$
 - for both i , $\sum_{s_i} \mathbf{p}_{s_i} = 1$
- \mathbf{b}_{s_i} is a binary variable indicating whether s_i is in the support, M is a large number